An Interop Publication

ConneXions®

The Interoperability Report

December 1992 Volume 6, No. 12

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.
E-mail: connexions@interop.com.

ISSN 0894–5926

From the Editor

Our report from INTEROP 92 Fall is going to be short and sweet: More than 55,000 attendees, 400+ exhibitors, and an INTEROPnet consisting of more than 80 miles of cable and some 4,000 interconnected devices. The number of attendees exceeded even our own optimistic estimates and is a tenfold increase since INTEROP 88.

During the show, a new *Internet Domain Survey* was released. This survey attempts to discover every host on the Internet by doing a complete search of the Domain Name System. In October there were 1,136,000 hosts in the Internet, an increase of 14.5% since July. INTEROP, the internetworking industry, and the Internet itself is growing at an incredible rate!

Our series *Components of OSI* continues this month with a look at *conformance testing*. The purpose of conformance testing is to increase the probability that different OSI implementations will be able to operate with each other. The article presents an overview of the OSI conformance testing methodology of ISO and CCITT and is written by Howard Motteler and Deepinder Sidhu of the University of Maryland.

We are nearly at the end of the OSI series and are discussing the possibility of compiling it all into a book, but first we want to make sure that everything is covered. If you have suggestions for other OSI articles, please drop us a line. A complete list of the articles to date is included on page 11.

In our October issue, we looked at network management from an architectural and standards perspective. This month, we switch focus to the users and bring you two network management case studies: one from the commercial sector and one from a university. The articles are by John Scoggin, Delmarva Power and Light, and Craig Finseth, University of Minnesota.

Our final article of the year is a look back to INET '92, the first annual conference of the Internet Society, which was held in June. Next year, INET '93 will be held in San Francisco immediately preceding INTEROP 93 Fall. For more information see page 32.

Two new Internet books, *The Internet Companion: A Beginner's Guide to Global Networking,* and *The Whole Internet User's Guide and Catalog,* are reviewed in our Book Review section. We also have a review of the book *XTP: The Xpress Transfer Protocol.* For more information about XTP, see *ConneXions,* Vol. 5, No. 12, December 1991.

With that, it is time to wish you happy holidays and welcome back to Volume 7 in a few weeks.

Components of OSI: Conformance Testing

by
Howard Motteler & Deepinder Sidhu, University of Maryland

Introduction

The purpose of conformance testing is to increase the probability that different OSI implementations will be able to operate with each other. This article presents an overview of the OSI conformance testing methodology of ISO and CCITT. We give ISO definitions of specification, conformance, and consider both abstract and concrete testing. (ISO acronyms are defined when first introduced, and we have also included an appendix of those acronyms used here.)

ISO conformance testing [1] is a general method for testing OSI products for conformance to relevant OSI standards [2–6]. Major components of ISO conformance testing are the specification of abstract test suites, the implementation of these test suites, and the production of test reports. Standardized test suites must be developed for each International Standard or CCITT Recommendation specifying an OSI protocol [1]. The standardized test suites for an OSI protocol can be used by its implementor for self-testing, by users for validation, and by third party testing organization for certification. ISO conformance testing is also applicable to conformance testing of a standard specifying transfer syntax (if implemented in conjunction with OSI protocols) and in principle to conformance testing of ISDN two-party protocols.

Limitations

There are some limitations in the scope of the ISO methodology. There is no testing by application-, protocol-, or system-specific methods, no testing of protocol-independent requirements, and no test methods involving more than two end-systems in communication. ISO conformance testing does not involve tests for performance or robustness of an implementation. Conformance testing can not pass judgement on any pure implementation decision, or on the way services are implemented or actually provided.

The complexity of real-world protocols makes exhaustive testing costly and impractical. The purpose of conformance testing is to increase the probability that different OSI implementations will be able to operate with each other. It provides confidence that an implementation has some required capabilities and behaves correctly in selected instances of its communication. Conformance is only required for the external behavior of an OSI implementation, even though a standard may define both internal and external behavior.

The ISO conformance testing methodology is defined in seven documents. (Five have been finalized, while work is progressing on two more.) The first five are: *General Concepts* [2], *Abstract Test Suite Specification* [3], *Tree and Tabular Combined Notation* (TTCN) [4], *Test Realization* [5], and *Requirements on test laboratories and clients for the conformance test assessment process* [6]. Figure 1 gives an overview of the relationship among these components.

Conformance in OSI

A real system exhibits conformance if it meets the requirements of applicable International Standards in its communication with other real systems. ("International Standard" here means OSI International Standards or CCITT Recommendations.) The system must conform to these standards both individually and as a group. Conformance requirements may be mandatory, may hold under specific conditions, or may be optional, and may be stated either positively or negatively. Conformance requirements may be *static* or *dynamic*.

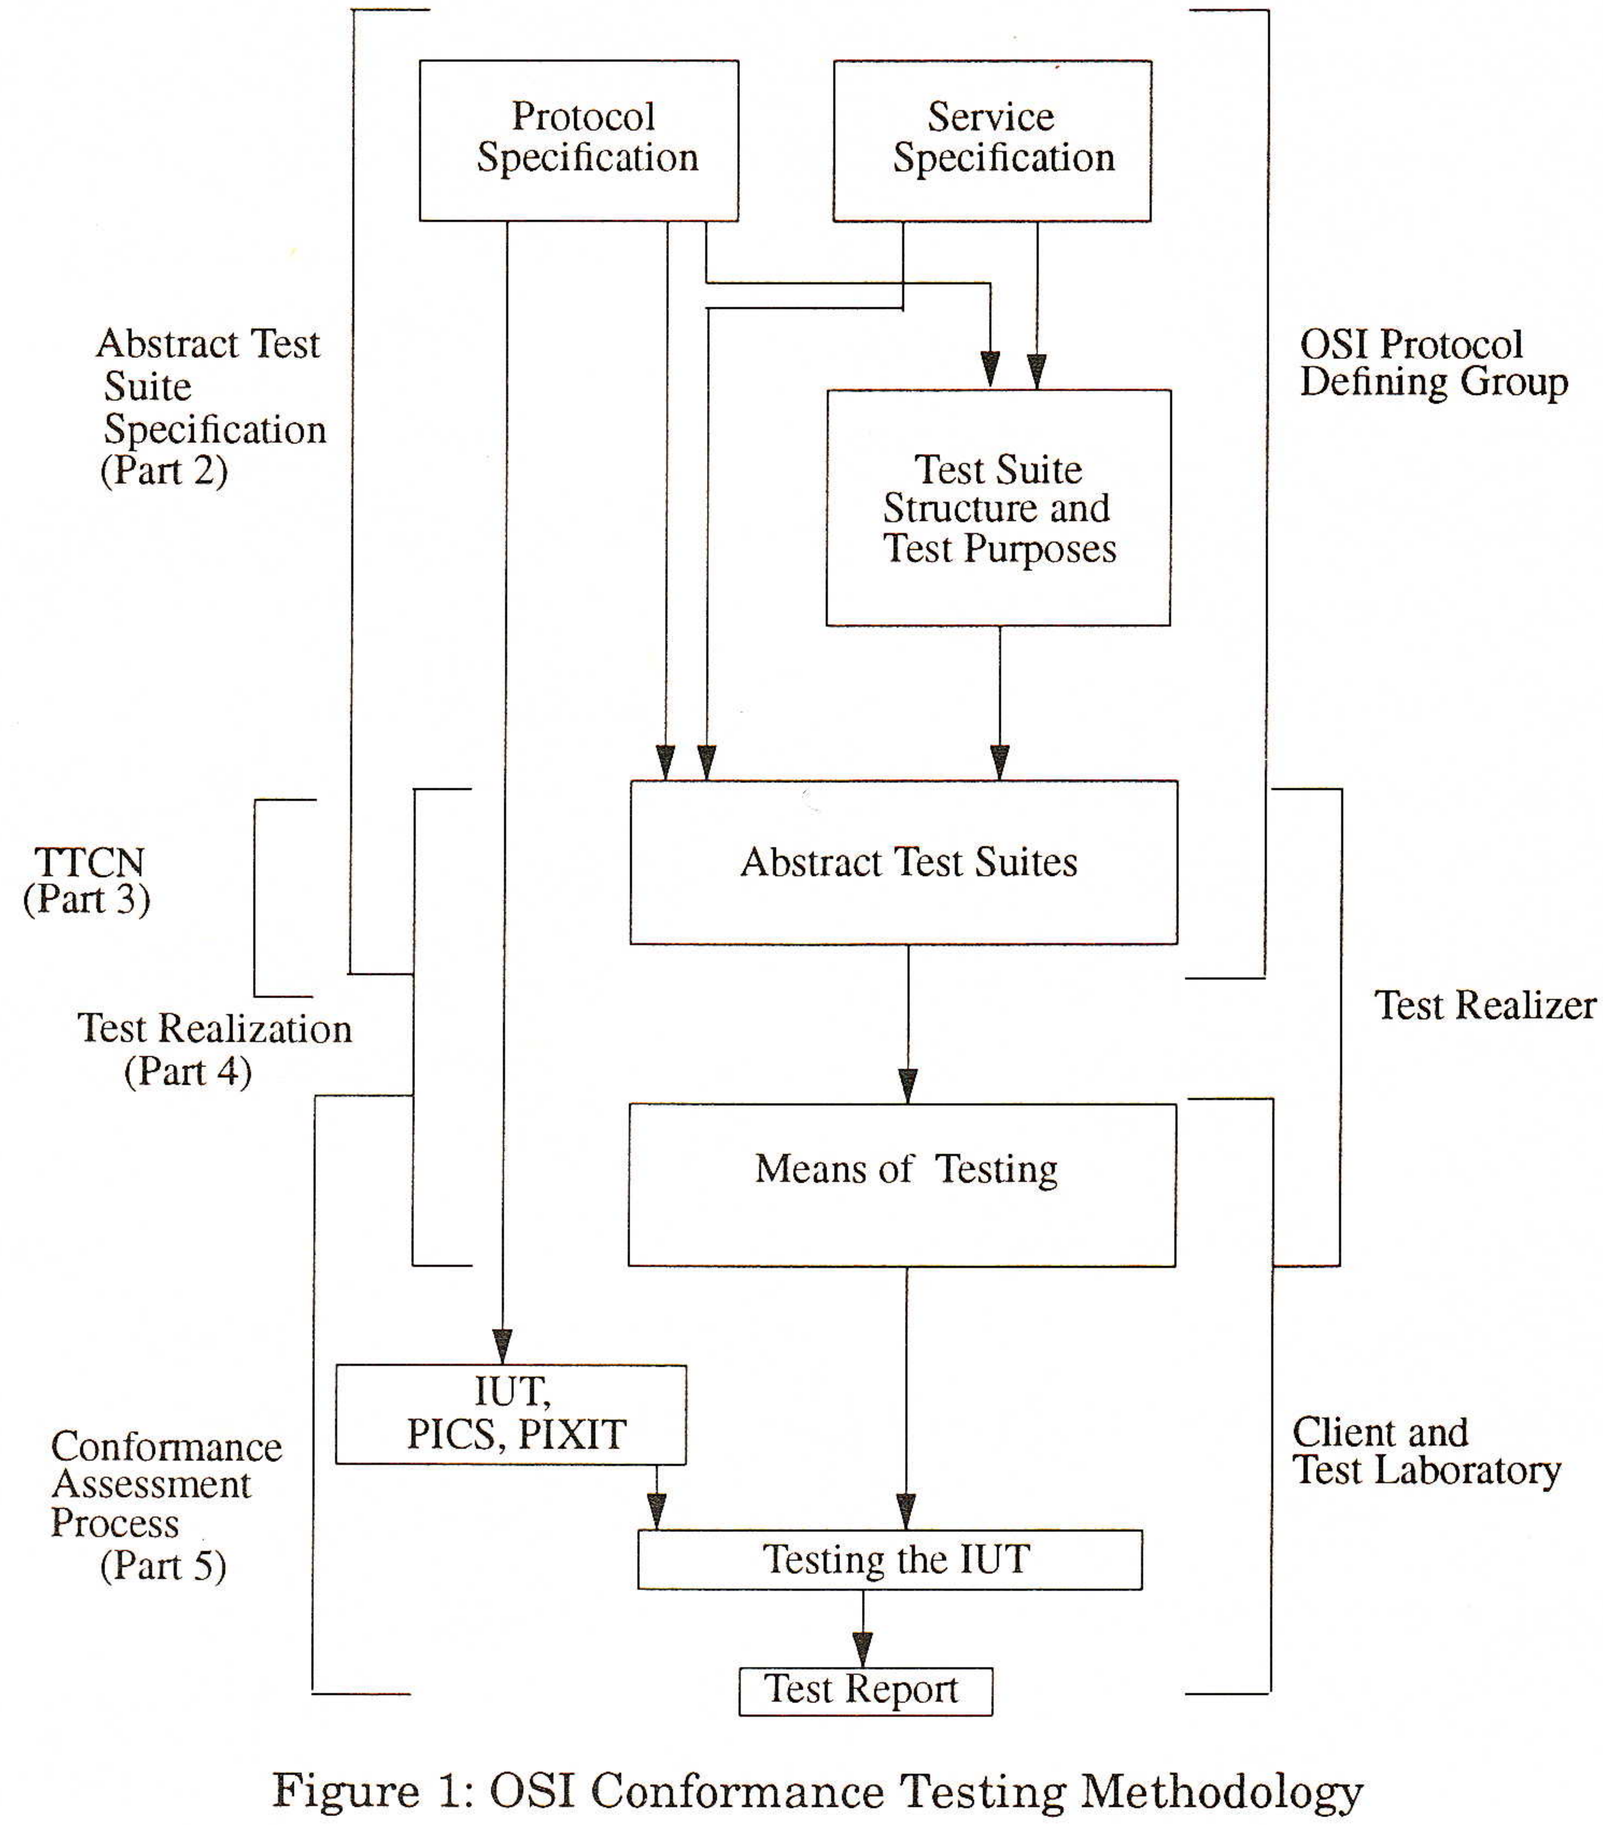

Figure 1: OSI Conformance Testing Methodology

Static conformance requirements specify limitations on the combinations of capabilities permitted in a real system, and define minimum capabilities for interworking. Static requirements might specify values taken by timers and other parameters. Static conformance requirements determine capabilities to be included in a protocol implementation, and may also determine multi-layer dependencies, by placing limitations on the capabilities of the underlying layers of the system in which the protocol implementation resides.

Dynamic conformance requirements specify observable communication behaviors. A system exhibits dynamic conformance in a communication instance if its behavior is a member of the set of all behaviors permitted by the relevant specification. Dynamic conformance requirements are those that define the actual protocol; the use and format of its *Protocol Data Units* (PDUs), state transitions, negotiation rules, etc.

PICS The *Protocol Implementation Conformance Statement* (PICS) is a statement of the capabilities and options that are present in a particular implementation. A separate PICS is produced for each protocol in a set of interrelated OSI protocols. PICS should make a distinction between the mandatory, optional, and conditional static conformance requirements of the protocol itself, and similar information related to multi-layer dependencies.

continued on next page

OSI Conformance Testing *(continued)*

In addition to PICS, a client submitting an implementation to a test laboratory provides information about the IUT (*implementation under test*) and its testing environment. This information is called *Protocol Implementation eXtra Information for Testing* (PIXIT). PIXIT contains information about the SUT (*system under test*), and information that adds specificity to information contained in PICS. It also helps in determining testable from non-testable capabilities.

The primary goal of conformance testing is to increase the probability that different implementations of an OSI standard are able to interwork with each other. Conformance is a necessary but not a sufficient condition for interworking. Interworking between two implementations is more likely to be successful if they conform to the same subset of a protocol; checking this requires a comparison of the PICS and system conformance statements. Different interpretations of ambiguities and selected optional capabilities in OSI International Standards often lead to implementations that are not able to interwork. Test reports from unsuccessful attempts at interworking and PIXIT provide additional information which help in making the two implementations compatible and hence assist with interworking.

Conformance testing types

As we have noted, exhaustive testing of OSI systems for conformance is impractical. Four types of testing have been identified to assess conformance of an implementation to a standard, as follows:

• *Basic Interconnection Tests* (BITs) provide very limited form of testing of an IUT with respect to the main features of the protocol. They check only the feasibility of interconnection with other IUTs. These tests are appropriate for checking test environment, deciding about further tests and determining interoperability with other conforming implementations. These tests alone are inappropriate for claiming conformance and detecting causes of failures.

• *Capability tests* allow checking of static conformance requirements stated in the PICS and validating observable capabilities. These tests are useful for checking consistency of PICS with IUT. They alone can not be used for claiming conformance or resolving problems encountered in actual testing.

• *Behavior tests* allow comprehensive checking of consistency with a full range of dynamic conformance requirements. These tests together with capability tests can form basis for conformance assessment of an IUT to its standard.

• *Conformance resolution tests* are non-standardized tests. They are used for testing parts of the protocol which can not be tested using a standardized *abstract test suite* (ATS). Such situations can arise, for example, due to limitations of the chosen test method, untestability of a conformance requirements or a new situation arising during implementation. Conformance resolution tests are performed using the diagnostic and debugging facilities of SUT or local operating systems.

Conformance testing results

The conformance assessment process is the complete set of activities involved in demonstrating conformance of an IUT to a standard. It consists of test preparation, operation, and report production. Test preparation involves production of the System Conformance Statement, PICS and PIXIT, choice of Abstract Test Method and Abstract Test Suit (ATS), and preparation of the SUT and Means of Testing.

Test operations involves analysis of PICS for compliance with static conformance requirements, selection and parametrization of test based on the PICS and PIXIT, and selection of one or more test campaigns. A test campaign is the entire activity of executing the *Parameterized Executable Test Suite* (PETS) and recording of observed sequences of test events in a conformance log. It consists of BITs (optional), capability tests, and behavior tests. Finally, test report production involves analysis of results and production of a report. The three phases of the conformance assessment process are shown in Figure 2.

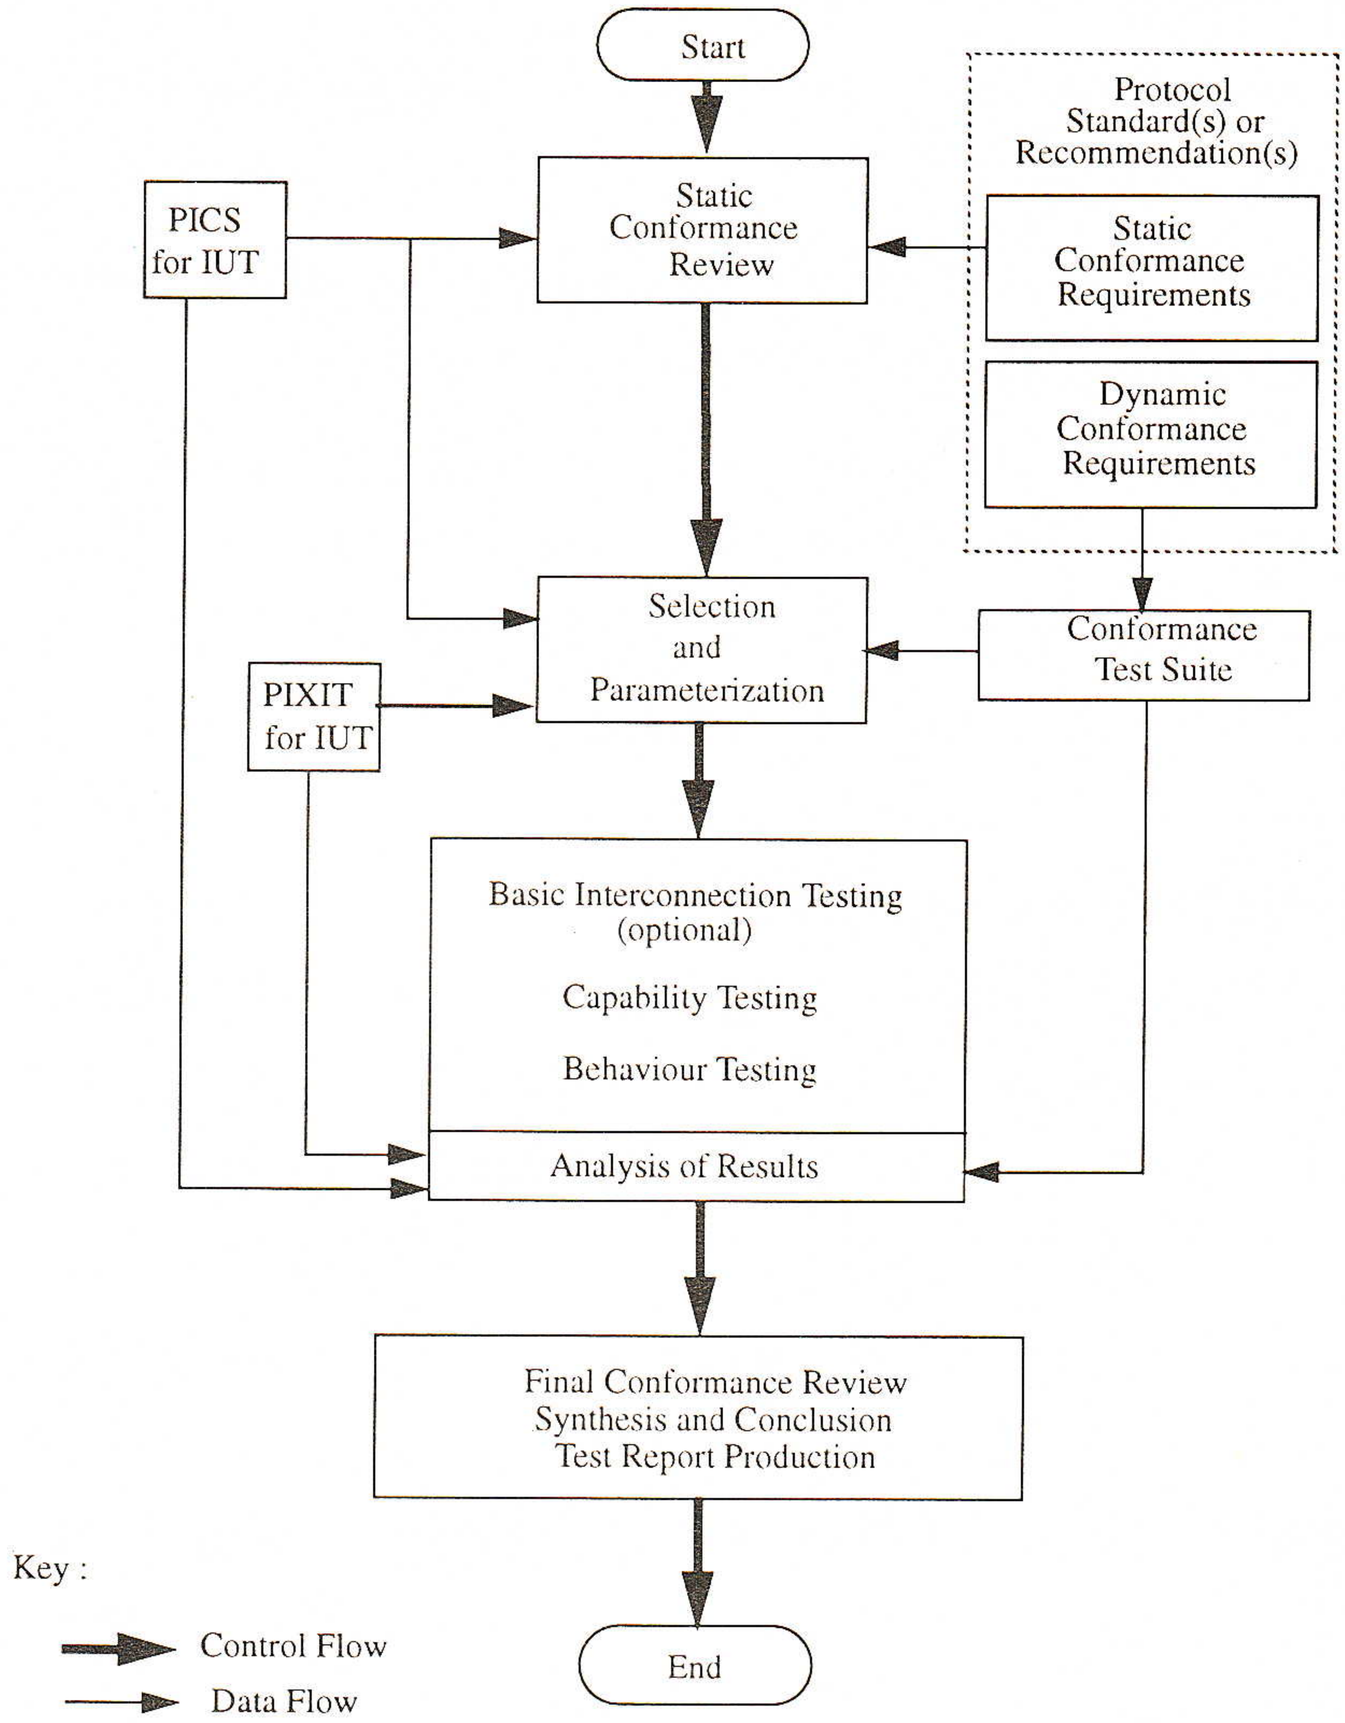

Figure 2: The Conformance Assessment Process

The outcome of a test is the sequence of events, i.e., inputs and outputs, that took place at the *point of control and observation* (PCO) of IUT. A test outcome can be either foreseen or unforeseen. An unforeseen outcome happens if events during execution do not match any sequence of events defined in the abstract test case. A foreseen test outcome is one which is defined in the abstract case, and may include unidentified test events. A foreseen test outcome results in one of three test verdicts: pass, fail, or inconclusive. The pass verdict means that the observed test outcome gives evidence of conformance to the conformance requirements. The fail verdict means that the observed test outcome either contains invalid test events or indicates nonconformance with respect to one or more of conformance requirements. An inconclusive verdict means neither a definite pass nor fail can be given.

continued on next page

OSI Conformance Testing *(continued)*

A credible conformance test must meet some basic requirements. The test results must be repeatable in that the results of executing a test case should be the same whenever it is performed. The standardized testing procedures applied to an IUT should produce comparable results whether they are performed by its implementor, a user or a third party. This requires careful specification of test cases, means of testing and procedures to be followed in repeating tests. The test procedure must log sufficient information so that auditing can be done to verify that correct procedure has been followed during testing.

The results of conformance testing are documented in two reports: *System Conformance Test Report* (SCTR) and *Protocol Conformance Test Report* (PCTR). The SCTR provides a summary of the conformance status of SUT and PCTR, one per protocol in SUT, documents the results of test cases with reference to conformance logs.

Conformance testing methods

Test methods vary according to configurations of real open systems to be tested. For example, test methods depend on the main functionality of the system (end-system or relay-system) and on the particular layers using OSI protocols. Three basic configurations defined for conformance testing are: (1) 7-layer open systems (end-systems), which use OSI standardized protocols in all seven layers, (2) Partial (N)-open systems (end-systems), which use OSI standardized protocols in layers 1 to N, and (3) Open relay-systems, which use OSI protocols in layers 1 to 3 (network relay-system) or 1 to 7 (application relay-systems). The basic configurations can be combined to define other configurations.

A system under test (SUT) may be a combination of basic configurations 1 and 2 with the choice of OSI and non-OSI protocols above layer N. An IUT is part of a real open system which implements one or more related OSI protocols in the same layer or adjacent layers. The IUTs can be single-protocol or multi-protocol implementations. If an IUT is a relay-system then it contains at least the layer containing relay function. An IUT defined as an open relay-system must include at least the layer which provides the relay function.

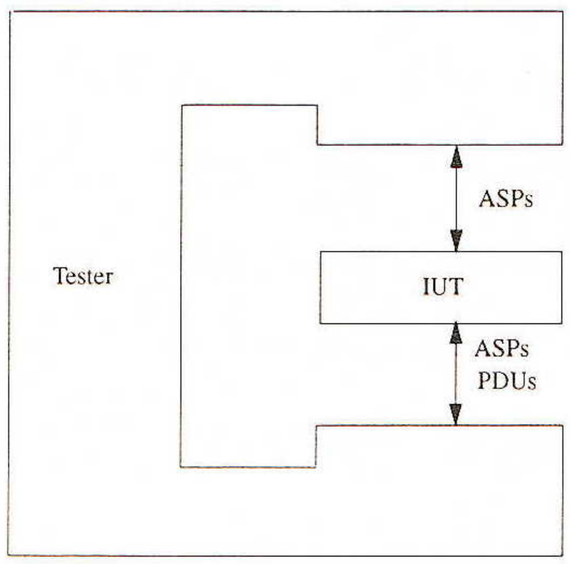

Figure 3: Conceptual Testing Architecture

Test methods are based on an abstract testing methodology, and use the OSI reference model. An abstract test method is specified by identifying the points, closest to the IUT, at which control and observation of the behavior of an IUT is done. (The allowed behaviors of an OSI protocol are specified in its standard document.) The behavior of an (N)-entity is defined in terms of the service primitives above and below it, i.e., (N)-ASPs and (N-1)-ASPs where the latter include (N)-PDUs. In a multi-protocol IUT, the behavior is defined in terms ASPs above and below it and also PDUs of the protocols in it. Figure 3 shows a conceptual tester, an IUT and its interactions.

PCO

The interactions can be observed and controlled in different ways. Each point of control and observation (PCO) is identified by the service interface, in the OSI model, at which the test events are controlled and observed, by the set of test events (ASPs or PDUs) that are controlled and observed at this point, and by whether test events are controlled and observed within the SUT or in the test system. The activity associated with ASP below the IUT can be observed and controlled by the peer activity in a test system via an underlying service provider. This service provider is assumed to be sufficiently reliable. The ASP activity above the IUT may not be controllable and observable.

Abstract test methods

The *abstract test method* is a form of "black-box" testing, based on specifications for the IUT. Abstract testing methods use a *Lower Tester,* an *Upper Tester,* and some sort of *Test Coordination Procedure.* The Lower Tester provides means, during test execution, of indirect control and observation of the lower service boundary of the IUT with the help of the underlying service-provider. The Upper Tester provides means, during test execution, of control and observation of the upper service boundary of the IUT as defined by the chosen Abstract Test Method. The Test Coordination Procedures provide rules for cooperation between the Upper Tester and Lower Tester. In some cases, the rules of cooperation can be formally defined as a *Test Management Protocol.*

Four types of abstract test methods have been defined for IUT in the end-system SUT. All four test methods use control and observation of ASPs below the IUT and PDUs exchanged with the IUT, by means of a Lower Tester separated from the SUT, and possibly control and observation of ASP above the IUT. The four abstract test methods are shown in Figure 4 where Nt and Nb refer to the highest and lowest numbered layer in the IUT. For a single layer IUT, Nt is equal to Nb.

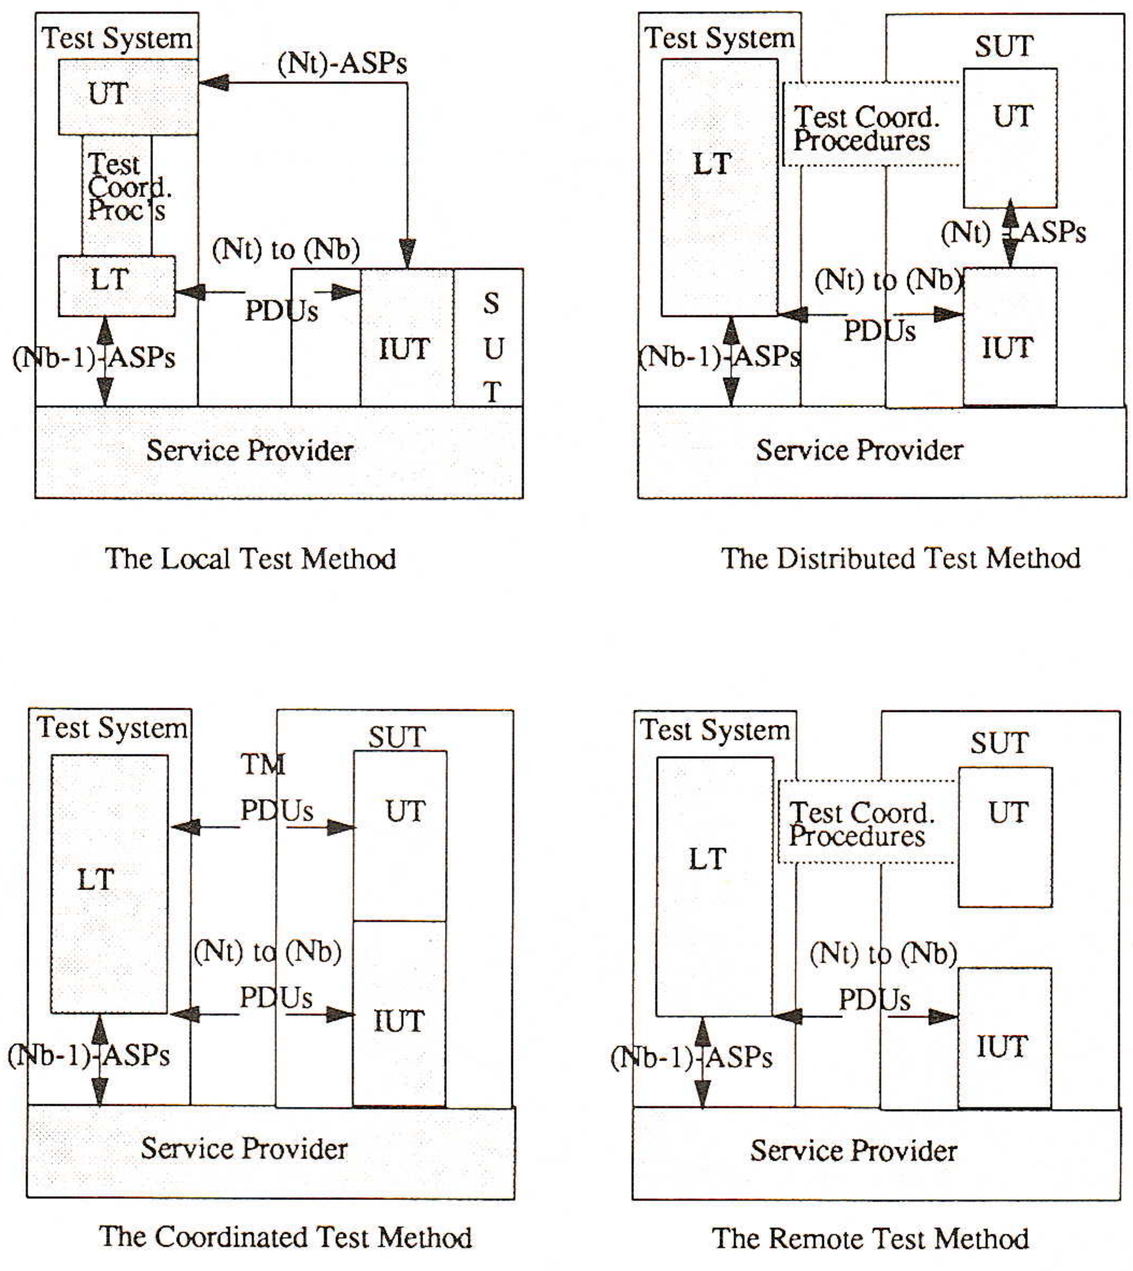

Figure 4: Abstract Test Methods

continued on next page

OSI Conformance Testing *(continued)*

The *Local Test Method* consists of two Points of Control and Observation (PCOs), one beneath the Lower Tester and the other at the upper service boundary of the IUT. The Upper Tester is located within the Test System. The upper service boundary of the IUT is a standardized hardware interface. The test coordination procedures are realized within the test system. Local test methods are applicable only to testing SUTs that have two hardware interfaces.

The *Distributed Test Method* consists of two PCOs, one beneath the Lower Tester and the other at the upper service boundary of the IUT. The Upper Tester is also located within the SUT. The upper service boundary of the IUT is either a human user interface or a standardized programming language interface. Distributed test methods are applicable only to testing IUTs that have an upper interface accessible either to a human user or to a software Upper Tester with a standardized programming language interface.

The *Coordinated Test Method* consists of only one PCO beneath the Lower Tester. It does not require access to the upper service boundary of the IUT. The test coordination procedures are realized by means of standardized *Test Management Protocols* (TMPs). The Upper Tester is an implementation of the TMP. Coordinated test methods are applicable where it is possible to implement a standardized TMP, in an Upper Tester in the SUT, above the IUT.

The *Remote Test Method* consists of only one PCO beneath the Lower Tester. It also does not require access to the upper service boundary of the IUT. The requirements for the test coordination procedures may be implied or informally expressed in the ATS, but no assumption is made regarding their feasibility or realization. There is no upper tester, some of its function are performed by the SUT. Remote test methods are applicable when it is possible to make use of some functions of the SUT to control the IUT during testing, instead of using a specific Upper Tester.

Single-layer test methods are appropriate for testing the majority of the protocol conformance requirements. Embedded test method variants permit the use of single-layer testing to all protocols of a multiprotocol IUT. For testing 7-layer open systems, the preferred test methods are the appropriate single-layer embedded test methods, used incrementally, with the PCOs: the upper interface of the Application layer as provided by the 7-layer open system, when applicable; successively, each SAP (or corresponding PCO if there is no SAP as such) below the protocol which is the focus of the testing, as controlled and observed in the Lower Tester, starting from the lowest protocol of the IUT and working upwards.

A method for abstract specification

The ISO standards for protocol specification are *Estelle* and *LOTOS;* we briefly consider Estelle in the context of testing. Estelle is based on a *finite state machine* (FSM), extended with variables. Transitions in Estelle occur when messages are received, and in certain instances without external input. Transitions, in conjunction with state information from variables, typically result in some action, which may include sending a message. Messages may be ASPs from the layer above, and ASPs and PDUs from the layer below.

In a given state, the IUT, considered as a process, may be waiting to receive an input from either the layer above or the layer below. Thus, even if the FSM itself is deterministic, the IUT may act as a nondeterministic process in a larger system.

One result of this is that even simple protocols can have complex behaviors as components in a larger system.

One form of abstract test for such an FSM-based specification is simply an input sequence (or possibly a set of input sequences) paired with expected outputs. (This will only test the "control" portion of a specification, i.e., behavior on state transitions independent of internal variable values.) Even if the number of internal states in an implementation is known (a questionable assumption in practice) an exhaustive search of FSM behaviors may be too costly. This leads to test methods based on checking edges; essentially, the idea is that if states can be identified (via some sort of signature) then edges can be reliably checked. The theory of this sort of testing has been extensively investigated [1].

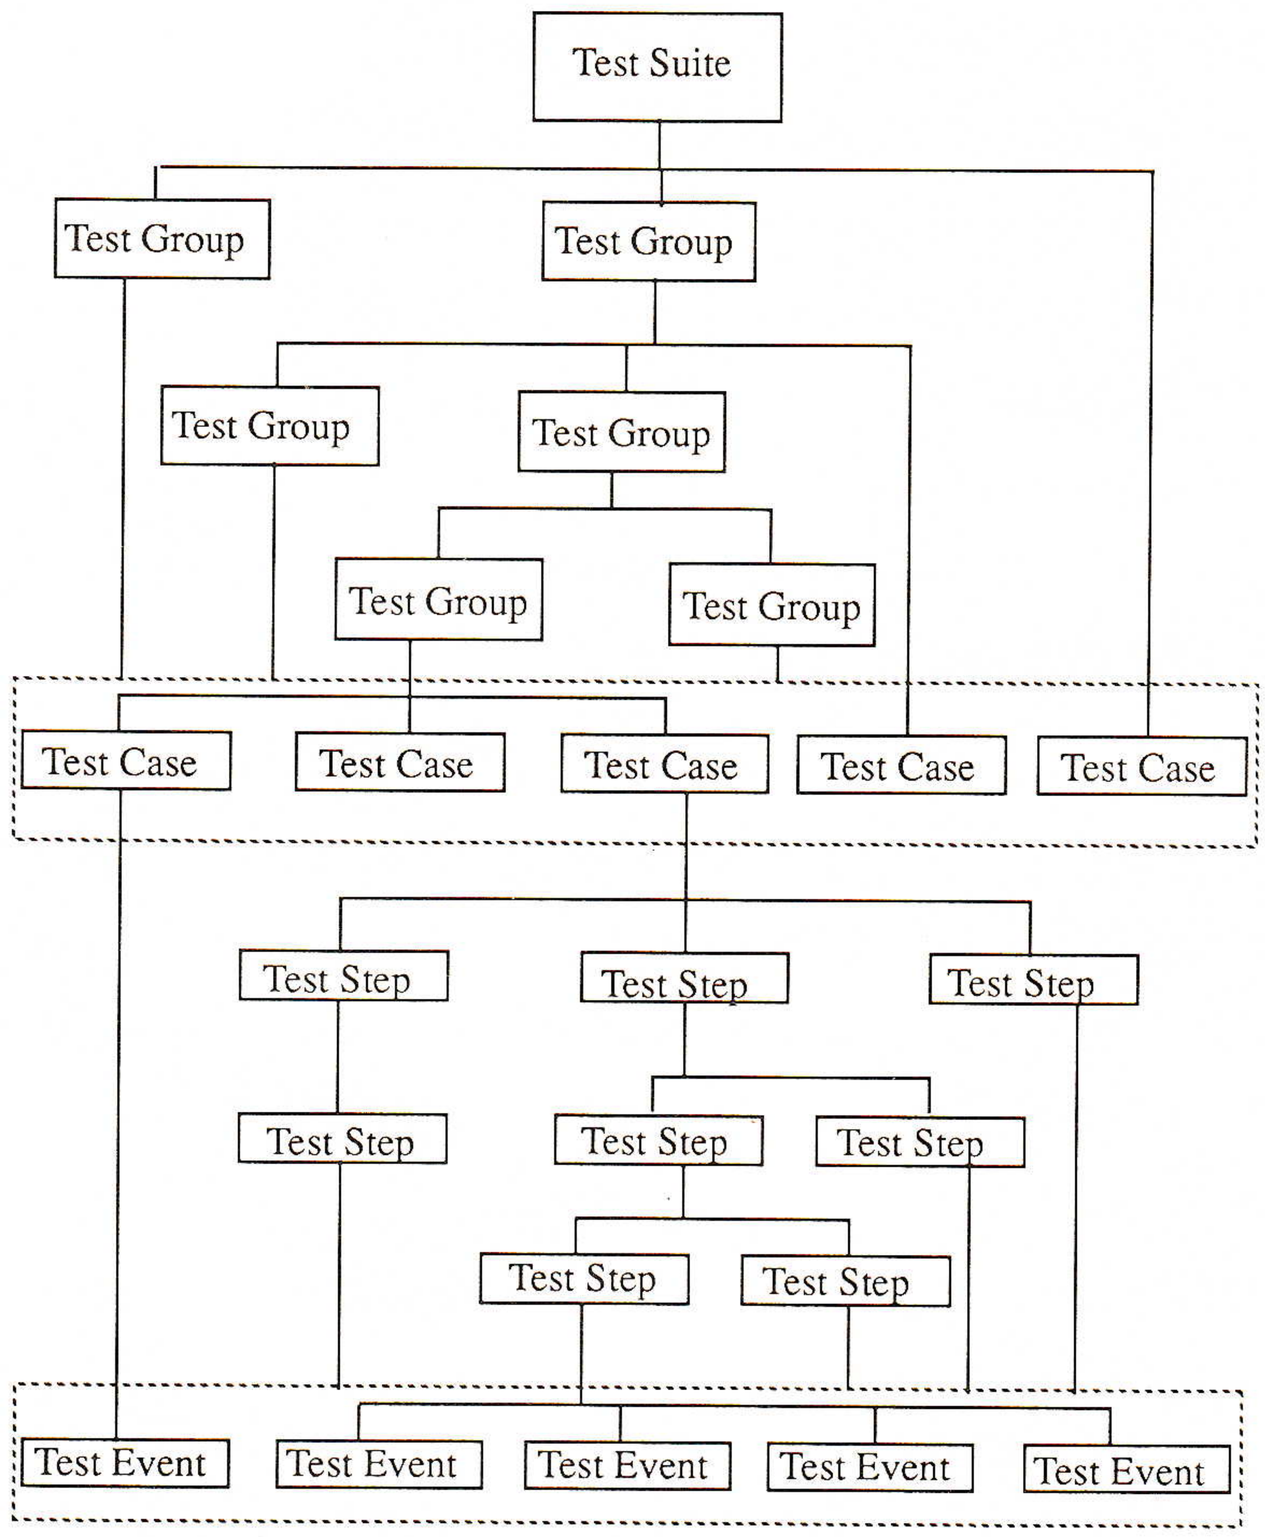

Figure 5: Test Suite Structure

Conformance test suite

Test suites have hierarchical structures as shown in Figure 5. A test suite is a complete set of test cases, possibly combined into nested test groups. Test groups provide logical ordering of test cases and aid in the development and execution of test suites. Each test case has a specific purpose such as testing a specific capability (e.g., support for a specific value for a parameter) of an IUT. Test cases can be refined into *test steps*. Common test steps can be grouped into nested *test step libraries*.

continued on next page

OSI Conformance Testing *(continued)*

A test step library can be associated with a test suite or with a specific test group in it. A *test event* is an indivisible unit of specification at the level of abstraction of the Spec., e.g., sending or receiving a PDU.

An abstract test case is derived from a test purpose or combination of test purposes. It specifies all sequences of foreseen test events which are necessary to achieve the test purpose. These test events constitute the test body. It also specifies a test preamble, which is a sequence of test events to put the IUT into the initial testing state for the test body, and test postamble, which is a sequence of test events to return the IUT to the desired stable testing state. A stable testing state of an IUT is one that can be maintained, independent of the lower tester, for time period longer than the gap between the end of one test case and beginning of the next. An executable test case is derived from an abstract test case.

Conclusion

ISO conformance testing plays a major role in guaranteeing interoperability. We have described major components of ISO conformance testing: definitions of specification and conformance, and both abstract and concrete testing methods.

References

[1] D. P. Sidhu, *OSI Conformance Testing,* Prentice-Hall, Englewood Cliffs, New Jersey, (to be published in 1993).

[2] ISO/IEC 9646-1, "Information Technology—Open Systems Interconnection—Conformance Testing Methodology and Framework—Part 1: General Concepts," 1991.

[3] ISO/IEC 9646-2, "Information Technology—Open Systems Interconnection—Conformance Testing Methodology and Framework—Part 2: Abstract Test Suite Specification," 1991.

[4] ISO/IEC 9646-3, "Information Technology—Open Systems Interconnection—Conformance Testing Methodology and Framework—Part 3: The Tree and Tabular Combined Notation (TTCN)," 1991.

[5] ISO/IEC 9646-4, "Information Technology—Open Systems Interconnection—Conformance Testing Methodology and Framework—Part 4: Test Realization," 1991.

[6] ISO/IEC 9646-5, "Information Technology—Open Systems Interconnection—Conformance Testing Methodology and Framework—Part 5: Requirements on Test Laboratories and Clients for Conformance Assessment Process," 1991.

Selected ISO acronyms

ASP	abstract service primitive
ATS	abstract test suite
BIT	basic interconnection test
FSM	finite state machine
ISDN	integrated services digital network
ISO	International Organization for Standardization
IUT	implementation under test
OSI	open systems interconnection
PCO	point of control and observation
PDU	protocol data unit
PETS	parameterized executable test suite
PICS	protocol implementation conformance statement
PIXIT	protocol implementation extra information for testing
SAP	service access point
SCTR	system conformance test report
SUT	system under test
TTCN	tree and tabular combined notation

HOWARD MOTTELER is an Assistant Professor of Computer Science at University of Maryland Baltimore County, currently on leave for the '92–'93 academic year at NASA/GSFC with an NRC research associate award. Research interests center on parallel and distributed systems, including task assignment, conformance testing, and connectionist processing. Current projects include further work in conformance testing and task assignment, and investigating applications of neural nets for retrieving atmospheric temperature and composition profiles from spectroscopic data. Howard Motteler received his B.S. (1980) from University of Puget Sound, his M.S. from Purdue (1982), and his Ph.D. from University of Maryland at College Park, (1987). E-mail: `motteler@umbc3.umbc.edu`.

DEEPINDER SIDHU received his B.S. degree in Electrical Engineering from the University of Kansas, and the M.S. and Ph.D. degrees in Computer Science and Theoretical Physics respectively from the State University of New York, Stony Brook. He worked at Rutgers University, Brookhaven National Laboratory, Mitre Corp., SDC-Burroughs Corp. (now Unisys), and Iowa State University. At SDC-Burroughs, he managed the Secure and Distributed Systems department within the R&D Division. He is currently Professor of Computer Science with the University of Maryland Baltimore County (UMBC) campus and the University of Maryland Institute for Advanced Computer Studies (UMIACS) at College Park. He has published over 100 papers in Theoretical Physics and Computer Science. His current research interests are in the areas of computer networks and distributed systems. He is the Editor-in-Chief of Journal of High Speed Networks, a new international journal. He is a co-founder of Protocol Development Corp. (acquired by Phoenix Techn. Ltd.) and TeleniX Corp. He is the author of a graduate level text, OSI Conformance Testing, (Prentice Hall, 1993). E-mail: `sidhu@umbc3.umbc.edu`.

"Components of OSI" in *ConneXions*

In an attempt to explain the *Open Systems Interconnection* (OSI) model, promulgated by the International Organization for Standardization (ISO) and the Consultative Committee for Telephone and Telegraph (CCITT), *ConneXions* has been running a number of articles under the heading "Components of OSI." Here is a list of them:

Integrated Services Digital Network (ISDN)	April	1989
X.400 Message Handling System	May	1989
X.500 Directory Services	June	1989
The Transport Layer	July	1989
Routing overview	August	1989
IS–IS Intra-Domain Routing	August	1989
ES–IS Routing	August	1989
The Session Service	September	1989
Connectionless Network Protocol (CLNP)	October	1989
The Presentation Layer	November	1989
A taxonomy of the players	December	1989
The Application Layer Structure	January	1990
File Transfer, Access, and Management (FTAM)	April	1990
The Security Architecture	August	1990
Group Communication	September	1990
X.25—the Network, Data Link, & Physical Layers	December	1990
The Virtual Terminal ASE	January	1991
Systems Management	April	1991
CO/CL Interworking	May	1991
Open/Office Document Architecture (ODA)	August	1991
Abstract Syntax Notation One (ASN.1)	January	1992
Broadband ISDN	April	1992
Synchronous Optical Network (SONET)	April	1992
Asynchronous Transfer Mode (ATM)	April	1992
Inter-Domain Routing Protocol (IDRP)	May	1992
The Remote Procedure Call (RPC) Service	June	1992
OSI Conformance Testing	December	1992
International Standardized Profiles	*Coming soon*	

Network Management of a Complex LAN/WAN Environment

by John K. Scoggin, Jr., Delmarva Power & Light Company

Network environment

Delmarva Power & Light Company owns and operates a private digital backbone of over 230 route-miles of fiber-optic and digital microwave systems. Racal-Datacom 9000 T-1 multiplexers provide bandwidth control for a variety of applications, including energy management, radio dispatch, corporate data communications and voice tie lines.

In 1989, Delmarva Power embarked on the construction of a multiprotocol internetwork which would interconnect IEEE 802.3 local area networks (LANs) in each office building and power plant. The *Delmarva InterNet* (DIN) consists of a core router network interconnected with multiple T-1 lines over Delmarva's private WAN facilities (Figure 1). Smaller offices are connected with high-speed data links (448–672 kilobits/second) using leased T-1 lines and the Racal 9000 multiplexer network. This high-speed network permits access to computing resources anywhere in the company with excellent response time and also permits incremental backups of remote network file servers at night from the Data Center. Protocols supported include TCP/IP, SNA, DECnet (Phase IV), and Xerox Network System (XNS).

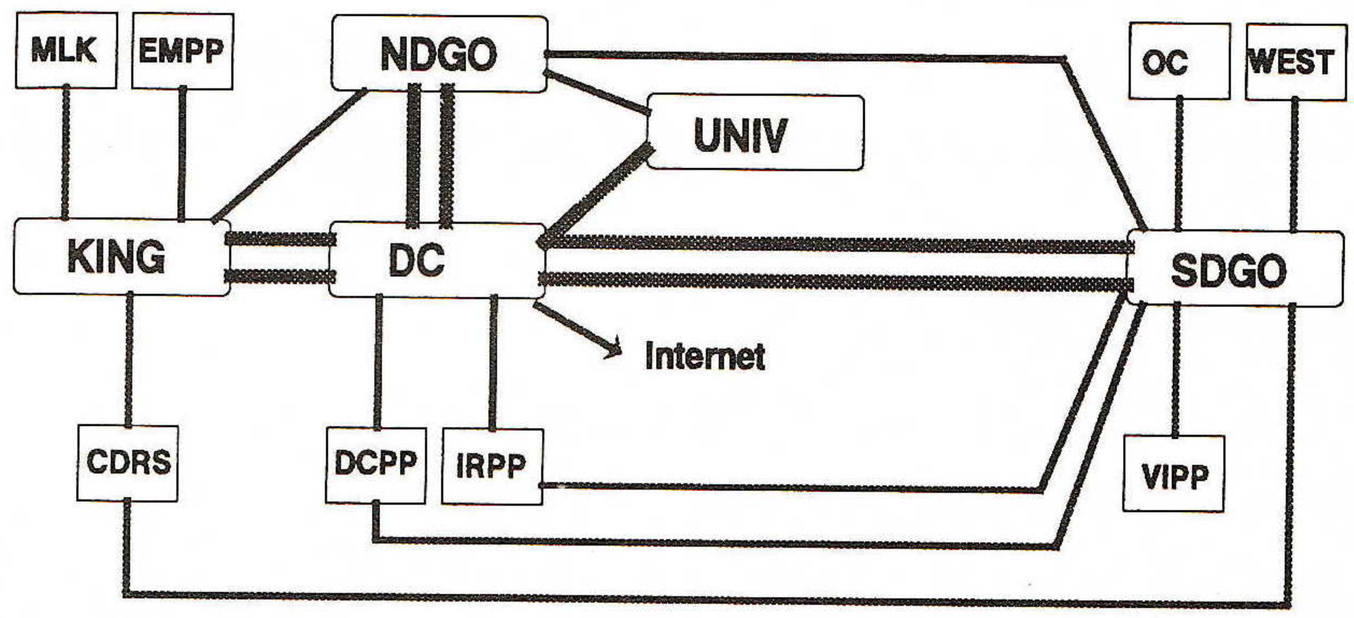

Figure 1: Delmarva InterNet — 1992

A wide variety of computing systems are attached to the Delmarva InterNet (DIN). An IBM 3090 provides central computing for administrative and engineering applications running in both MVS/ESA and VM/XA environments. A 3-processor complex of Control Data Cyber 180 systems is utilized to perform energy management and control. A number of mini-computer systems from IBM, Hewlett-Packard, Tandem, MIPS, Sun Microsystems, and others are utilized for a number of engineering and real-time control applications. All of these systems utilize the DIN for terminal access and data sharing.

The personal computer community utilizes the Banyan VINES network operating system. Approximately 600 machines are connected to thirty servers over the DIN, providing printer sharing, electronic mail, file systems, and communications services. TCP/IP is utilized for server-to-server communications.

Network Operations Center

The Network Operations section consists of ten full-time staff members and a variable number of temporary personnel and co-op students. Network Operations is responsible for the operation and maintenance of voice and data networks throughout the Corporation and performs all data network and distributed processing design and installation.

Due to the small size of the Network Operations staff, Delmarva has been involved in an aggressive automation effort since 1980. It is crucial that problems be identified and resolved in an efficient fashion. The *Network Operations Center* (NOC) located in Newark, Delaware performs all network monitoring, problem determination, and dispatch functions for telecommunications problems. They also function as the User Help Desk, providing first level support for all Information Systems-supported systems and equipment.

Network management architecture

The Delmarva Power NOC is equipped with a mixture of proprietary and standards-based network management systems. All major DIN components are compliant with the *Simple Network Management Protocol* standards. Wide Area Network components such as modems and multiplexers are managed using proprietary protocols, primarily due to their age and the lack of viable standards-based alternatives.

The NOC is equipped with a separate Network Management Network and a number of 80386/486 personal computers, VINES servers, and Sun SPARCstations. The Sun SPARCstation is the standard network management workstation in our NOC due to the large number of available applications and public domain code. An NFS (Network File System) server is used to act as a repository for all executables and databases.

Wide area management

Since Delmarva Power has standardized upon Racal-Datacom as its primary WAN vendor, we use several of their NMS systems as the basis for our operations. Low-speed network components (< 56 Kbps) are managed using the CMS400 network management system. This system is based upon an IBM PS/2 hub with remote PC workstations and provides fault detection and diagnostics for a variety of modems and multiplexers.

The T-1 transmission network is based upon Omnimux 9000 multiplexers that are managed by a Sun-based management system. This 9000 NMS provides a real-time graphical display of network status and permits bandwidth management.

Both of these network management systems are connected to a Racal CMS6000 Network Management System. This system is a "manager-of-managers"; alarms and status information from the CMS400 and CMS9000 systems are filtered and displayed in a common topology diagram depicting the entire Wide Area Network (Figure 2). Windows into the CMS400 and CMS9000 are also available. We are in the process of converting our T-1 Channel Service Units into "intelligent" units that are also managed by the CMS6000 product.

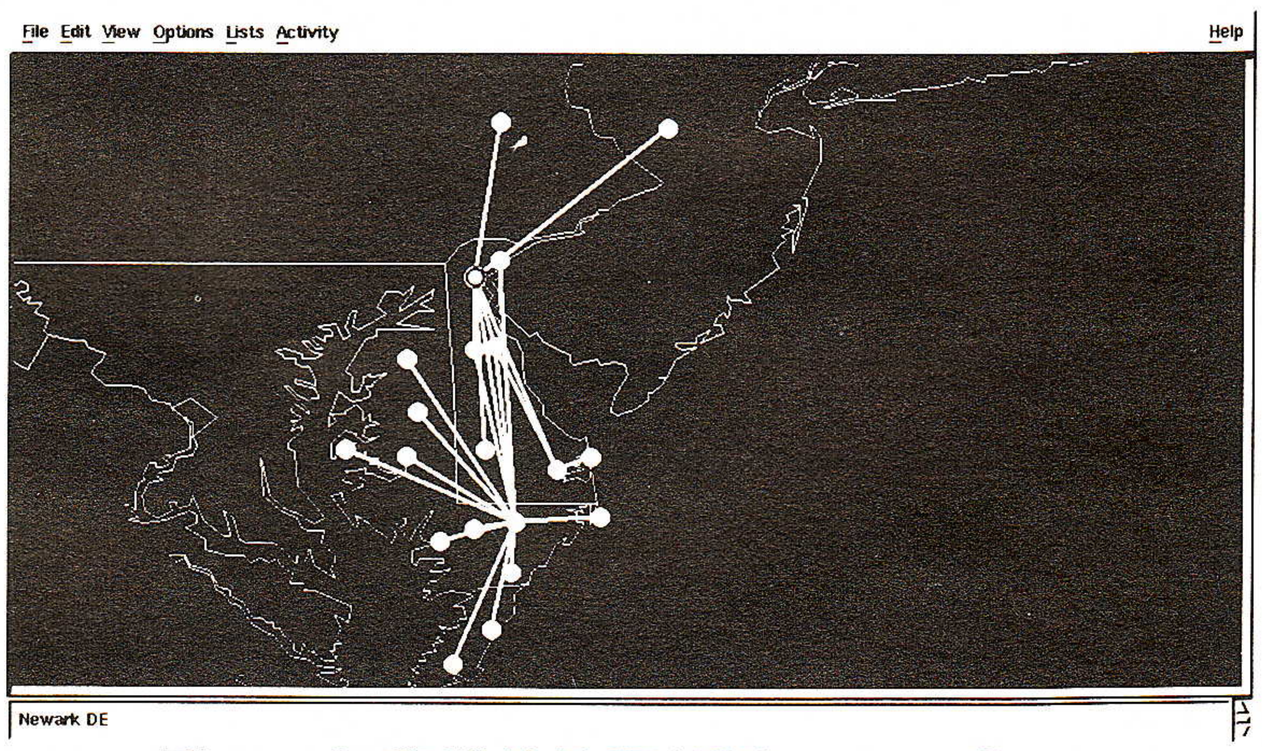

Figure 2: CMS6000 WAN Overview Screen

continued on next page

Management of a Complex LAN/WAN *(continued)*

CMS6000 is based upon client–server technology, with most user interface processing being done in the user's workstation. A Sun 4/75 provides central services to a number of remote Sun user workstations.

Delmarva also has several other "legacy" systems, such as a facility alarm system (Pulsecom DataLok-10) and Penril's 6000 network management system. These will be accessible from the CMS6000 as terminal emulators in the near future; alarm integration awaits the availability of an appropriate toolkit from Racal.

Delmarva InterNet management

The DIN was designed from inception as an SNMP-manageable system. Our routers (Wellfleet), wiring hubs (Cabletron), terminal servers (3Com), and file servers (Banyan, Intergraph, Sun) are all SNMP-compliant.

The Cabletron SPECTRUM product is used as our primary DIN management tool. SPECTRUM provides excellent alarm filtering and display for our entire installed base of equipment (Figure 3). SPECTRUM has proven quite scalable because its database/monitoring component is separable from the user interface component. A Sun 4/75 is utilized as the SPECTRUM server with five Sun 4/40 and 4/50 user workstations in concurrent operation.

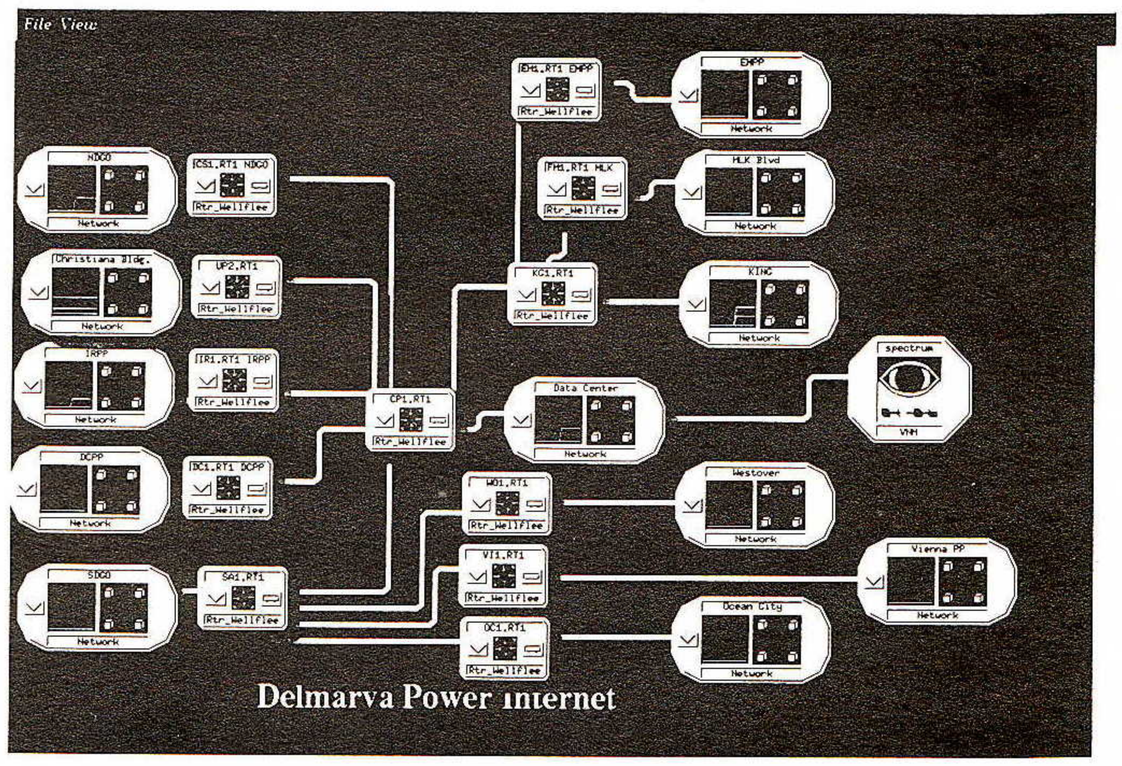

Figure 3: Cabletron SPECTRUM Topology Diagram

Help Desk operations

Delmarva utilizes an internally developed Help Desk Management System (HDMS) based upon the Unify Corporation's Unify 2000 DBMS and ACCELL 4th generation language. HDMS is used for inventory management, problem and project tracking, spares management, and budget preparation. HDMS currently runs on a Sun 4/75 system and is accessed using PCs, terminals, and Sun workstations.

HDMS is currently a character-based interface, but an X/Motif conversion will occur in the near future. Image support will be added at that time to simplify access to CAD drawings of important network systems.

Integrated environment

All of these network management systems operate in our Sun-based workstation network. A single IPC running in an X/Motif environment is capable of displaying complete network status using SPECTRUM and CMS6000 windows (Figure 4). Xterm sessions allow access to HDMS and other applications. Telnet 3270 is used to access the IBM 3090 environment through a McData 6100 TCP/IP gateway. This same workstation is also used for electronic mail and USENET news access.

Work is underway to permit access to several MS/DOS-based applications through the use of Desqview/X. Mail is gatewayed from the MS/DOS users to the Network Management Network using Banyan's SMTP Gateway software.

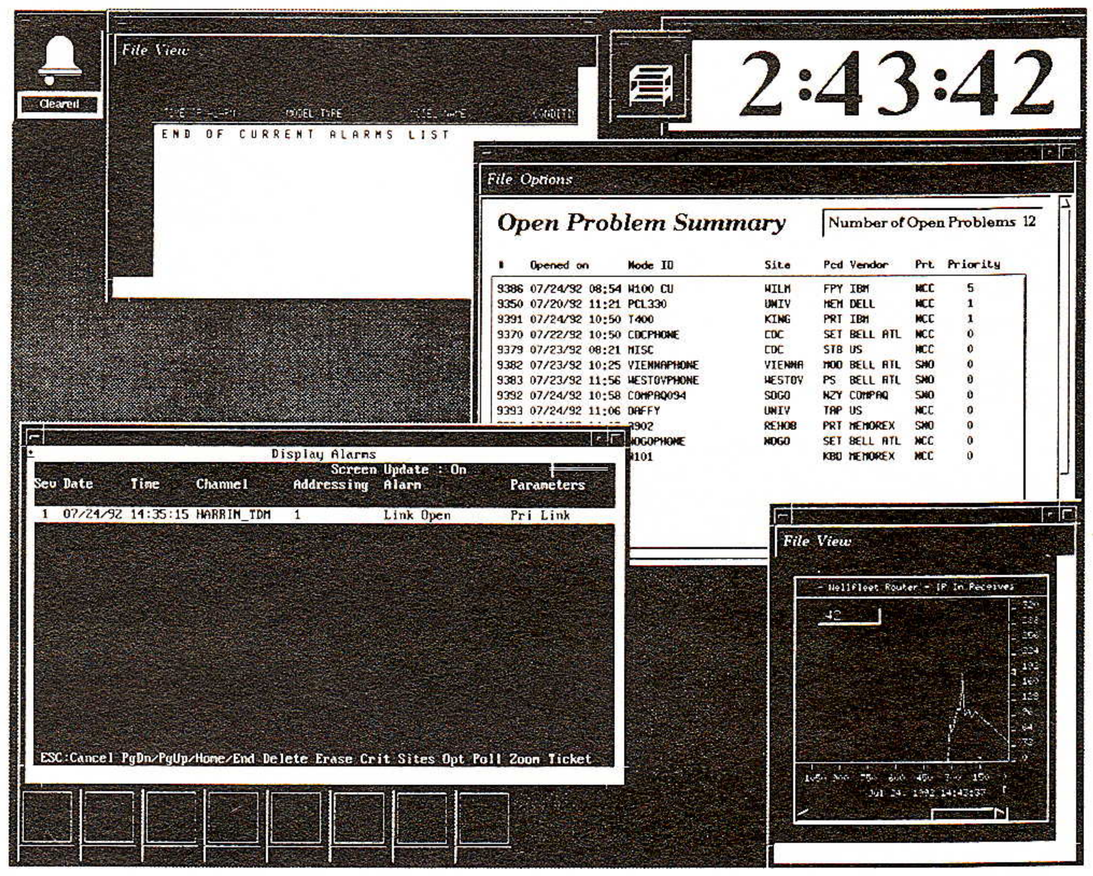

Figure 4: Delmarva Power Integrated NMS

Future plans

The future of IBM's *NetView* in our environment remains cloudy. We currently use NetView strictly as an SNA element manager. A NetView alarm parser has been written and will permit these events to be viewed in real-time using an X/Motif application. In the future, transfer of SNA network status to SPECTRUM would be highly desirable as we convert our entire SNA transport network to DIN. The role of SNA (at least in its current form) in Delmarva's network will diminish over time as we implement additional client-server applications.

Additional alarm parsers are in development for several "legacy systems," such as message switches and facility alarm systems. All of these parsers currently feed a standalone alarm management system (DIALS) based upon X/Motif. This information should be fed into the CMS6000 product whenever the appropriate APIs become available.

The implementation of the DIN has led to some interesting "pathological" TCP/IP routing problems in our Banyan VINES environment. Work is currently underway under a grant to the University of Delaware to develop a system to recognize these problems and display the condition is an easily understood fashion.

Conclusion

The level of integration attainable between network management systems from multiple vendors is largely limited to the sharing of user workstations at this time. The widespread usage of the X/Motif interface simplifies network operator training due its common "look and feel." In the longer term, data integration between these diverse management systems may become possible, possibly through the use of the more powerful manager-to-manager capabilities in the SMP protocol. The use of a common topology database server would simplify the maintenance of this information across multiple NMSs.

[Ed.: See also, "The Multi-Protocol Internet at Delmarva Power—A Case Study," in *ConneXions,* Volume 5, No. 10, October 1991.]

JOHN K. SCOGGIN, Jr. has been active in the data processing and telecommunications arena for 19 years. His current job responsibilities include the design, operation, and maintenance of a regional voice/data network and the operation of the User Help Desk. John is currently an Adjunct Assistant Professor in the Computer and Information Sciences Department of Goldey-Beacom College, teaching courses in Operating Systems and Data Communications. He can be reached as: `scoggin@delmarva.com`.

Thoughts on Network Management at the University of Minnesota

by Craig A. Finseth, University of Minnesota

Introduction

I recently attended a vendor presentation that described a new network management product. The product used advanced database technology and can store dozens of attributes on thousands of objects. Using this data, the product could automatically display many different views of the data in full color. It was customizable, extensible and in every way a state-of-the-art system. Yet I left the presentation feeling vaguely uncomfortable with what I had seen.

Consider this parallel: in the late 1970s, most U.S. Corporations believed that inventory was good and more inventory was better. They constructed huge warehouses and elaborate inventory tracking systems. Then came the 1980s and the realization that this inventory, rather then helping its owners, was actively hurting them. There were some major ancillary reasons for this: the inventory tied up working capital, its very maintenance and tracking drew resources away from other projects, and so on. Still, the real damage was that it masked over fundamental problems with the way that the corporation was organized. Lowering the level of inventory uncovered these problems, sometimes painfully. Correcting the problems resulted in major improvements in corporate operations and ultimately increased profits.

Functions

A network management system performs many functions:

- It tracks what equipment is where
- It tracks how the equipment is organized into a network
- It looks for malfunctioning equipment
- It provides information required to diagnose problems

So, if this parallel is to hold, what problems are uncovered by looking at information as the inventory of a network management system?

- *Ancillary:* I estimate that it would cost over $500,000 to create a database that described our existing network in the level of detail that this network management system supports.
- *Ancillary:* I estimate that it would cost over $200,000 each year to maintain the database.
- *Real Issue:* None of this investment by itself delivers any service to the users. None of this investment affects the mean time between failures, although it might have a tiny impact on the mean time to repair. (The reason for the tiny impact is simple: it doesn't matter if it takes one minute or one day to diagnose a problem if six weeks are needed to order and install the replacement unit.)
- *Real Issue:* Whenever incorrect data is encountered, it can significantly delay problem diagnosis. The incorrect data divert attention away from the real cause or to a spurious cause.

New directions

I suggest that these problems can be addressed by following these principles:

- *Assertion:* Data, left to itself, will deteriorate.
- *Principle:* Continually use the data (i.e., draw conclusions from it) and cross-check those conclusions to identify incorrect data.

- *Assertion:* If a datum is entered more than once, one instance will be wrong.
- *Principle:* Distinguish between entering data—that should only be done once—and referring to the data, which will be done often.
- *Assertion:* Creating a model of the network constitutes entering the data more than once.
- *Principle:* Have the actual network be its own model.

Data types

A network management system can be constructed using these principles that has four types of data: existence, sampled, entered, and derived.

Existence data are data implied by the existing tangible network. An example of such data is "Cisco IGS/L router, serial number 0002LP, located in the SW corner of room 4 of Pillsbury Hall." You obtain data of this type by traveling to the physical location(s) and directly observing the data. By definition, this type of data cannot be incorrect or out-of-date in any way.

Sampled data are obtained periodically (usually automatically) and stored for later reference. Such data are never completely up to date, since when they are transferred and stored, they may have changed. Often sampled data are somewhat static and the stored versions still may be useful for other purposes. Barring bugs in the implementation or transmission errors, these data are never incorrect. It is often useful to retain multiple versions of these data for comparison or analyses. Examples of this type of data are router configurations, routing tables, ARP tables, etc.

Entered data are (usually manually) entered and stored for later use. Typically, these data contain only the minimum information required to tie the other types of data together and are thus small in comparison. Examples of this type of data are DNS information, lists of routers, etc. (There are no hard boundaries between sampled and entered data. For example, a router configuration could be considered entered data.)

Derived data are created from other data. One example of this type of data is comparing sampled routing tables to entered lists of assigned network numbers and looking for discrepancies. Another example is expanding an entered list of devices by looking up addressing, maintainer, and other information and collecting that expanded information into one place.

In operation

Our current network management system is based on the above principles. While not completely implemented, enough is in place to demonstrate that this system is feasible.

We are currently making extensive use of sampled data. Each night, all routers are queried and the data saved. Each week, all Shiva *FastPaths* are queried.

Our entered data break down as follows:

1 MByte	Domain Name Server (DNS)
200 Kbytes	monitoring program configuration
200 Kbytes	other (network no. lists, contact information, etc.)

It is interesting that this amount of data can comfortably fit in current palmtop computers such as the HP95LX, with room left over for programs to access the data efficiently.

continued on next page

Thoughts on Network Management *(continued)*

We use the DNS to record the host name, IP address(es), and MX information. In addition, we optionally record the host and operating system types (very generic), the physical location if available, and the device's maintainer if different from that of the department. As comments, we record directives to a program that generates the bulk of the monitoring program configuration.

Our network contains about 10,000 hosts. The entered data thus averages about 140 bytes per host. (Much of this large size is the result of inefficient coding in the DNS.)

We are now using derived data to manage the *AppleTalk* network. We are also about to expand this type of data.

But wait, there's more!

If we were to stop here, our system would be failure. Like a flying buttress with no cathedral to lean on, it would fall over and we would have chaos.

The design of our network management system follows from the design of our actual network. To return to the original thesis: if the network itself is muddled, the network management system will reflect that mess by becoming complex. While one can improve the network management system to make the messy network survive, it cannot fix the disorder itself.

The purpose of a network is reliably to move data from one node to another. A network failure happens when the network does not fulfill its purpose. Failures are measured from the users' perspective. Failures are traditionally characterized by two parameters:

- *The Mean Time Between Failures* (MTBF) measures how often a failure is observed.
- *The Mean Time To Repair* (MTTR) measures how long it takes to rectify the failure.

We have selected the following values:

- MTBF of 1 year
- MTTR of 2 hours

These values were selected because we believe that they correspond to our users' desires. This is somewhat of an educated guess during our current network transition phase at the University. If asked, we believe that most of our users would *say* that they do not rely on the network and that it is a luxury. As such, it wouldn't matter if it went down or for how long it was down. Yet, we have not had a major network failure in recent memory. We believe that if we were to have such a failure, our users—and we—would quickly find out that they do not consider the network a luxury. A large-scale network failure probably will have the same consequences as a large-scale power failure (i.e., bring operations to a halt).

The following sections will describe the techniques we use to meet these design goals.

Defining the service

If you can't define your service, you can't measure its reliability. Our service is the "reliable," end-to-end delivery of packets from the originating node to the recipient. The hardware level interface is Ethernet or IEEE 802.3, at either an AUI or 10BaseT interface. The software level interface is any of TCP/IP, DECNET, AppleTalk Phase II or Novell IPX, with plans to add ISO/OSI in the future.

Reliable is in quotes in the above definition because it is not used to mean "100%," but adequate reliability as required by each protocol. For TCP/IP, for example, even a 2% or so failure rate at the packet level does not interfere with delivering reliable service to the user.

Keeping failures from happening

Failures can occur anywhere in the network. There are two ways to reduce the number of failures: either reduce the number of network components or reduce the failure rate of each component.

Given the typical size of buildings, space between buildings, and the number of network nodes, there isn't too much that can be done about the number of components. A single, physical network can only reach a tiny fraction of the existing nodes. Thus, multiple physical networks must be joined into a larger network by active components.

The next step is to reduce the effective failure rate of each component. This reduction is obtained by selecting reliable components and minimizing stress on each component.

For active components, we look at the observed failure rate and only choose components that are reliable. We minimize stress on these components by careful installation, minimal disturbance (e.g., locked rooms), limiting environmental stress (e.g., air conditioning), and using UPS power.

For passive components, we install them carefully, in full accordance with network specifications and use conservative network designs. For example, keeping network segments short and having only a few devices (preferably two) on each network segment (e.g., one host per twisted pair hub port).

Network control is another area of concern. Where possible, the active devices (e.g., routers) are configured only to accept control information from each other or a secure command area. The computers that hold the network management data are configured to be secure and all changes in network management data must be made from those systems. (Network information is available on a read-only basis to many other sites.)

Example

We will close this section by reviewing a typical cross-section of the network between a user and the server that they are using:

- User's machine (Macintosh)
- LocalTalk network
- Shiva FastPath
- Ethernet network
- Bridge
- Ethernet link segment
- 10BaseT hub
- Ethernet link segment
- Cisco IGS/L router
- Fibre Ethernet link segment
- Cisco AGS+ router
- FDDI ring
- Cisco AGS+ router
- Fibre Ethernet link segment
- Cisco IGS/L router
- Ethernet network
- Server host (mainframe computer)

continued on next page

Thoughts on Network Management *(continued)*

There are 15 network elements between the user's computer and the server. (We are not responsible for the reliability of any of the end nodes.) If the user is to observe a failure rate e of one per year on this network, each component must fail no more often than:

$$1/\text{year} \geq \{ 1 - (1 - e)^{15} \}$$

or about once in fifteen years. Now, fifteen years are about 130,000 hours and active equipment usually is rated at no more than 50,000 hour MTBF. Thus, if the passive elements (cabling) are considerably more reliable than necessary, we just might make our goal.

Fixing failures quickly

No matter how carefully a network is designed and installed, it *will* still fail. Therefore, it is important that any failures are repaired quickly. How quickly? With 10,000 nodes and a MTTR goal of two hours, the network has an annual failure budget of 20,000 node-hours. This is a large enough value to allow us leeway to solve a few tough problems—so long as most problems are solved quickly.

One thing we can't afford is an intermittent failure. Such failures are difficult to diagnose and resolve. It is not unusual for one to take a week to solve. If only ten nodes are affected, a single failure would eat up over eight percent of our failure budget.

The way to avoid intermittent failures is in network design. One characteristic of an out-of-specification network is its failure mode. Such a network will appear to work perfectly under low load conditions. Yet, once network load exceeds a threshold value, the network fails. As this threshold can be exceeded for periods roughly milliseconds at a time, the resulting effect is one of sporadic, inexplicable failures. This threshold value varies depending upon the exact way in which the network fails to meet specifications. The only way to avoid this class of problem is by strict adherence to network standards.

Still, even network standards make assumptions, and these assumptions may not be true. Amplifiers and receivers age and lose power, cable joints may be imperfect, cable lengths and transceiver spacing (for thick and thin net) may not be correct, interference may be present, network interfaces may not follow specifications, and so forth. For these reasons, we design as follows:

- If the network medium cannot be directly monitored (e.g., thick and thin net), stay well under network standards. For example, keep thin net segments to 100 m.
- If the network medium can be directly monitored, specifications may be met exactly or even exceeded slightly. This will only occur on 10BaseT and fibre Ethernet link segments where we have SNMP-compliant equipment on *both* ends of the link. With this configuration, we can (and do) directly track the error rate of the link and can tell if we have current problems or project future problems on the link. With intermittent failures minimized, regular failures can still cause problems. Given that the failure has happened, we:
- Keep the scope of any failure as local as possible. Our initial goal is that any failure should affect at most a single building. This constraint helps in two ways:

1) Fewer people are affected, therefore our failure budget is used up more slowly.

2) Limiting the scope also limits the possible causes, thus allowing us to locate the failure more quickly.

- Use only equipment that is "too simple to fail" (e.g., a cable) or smart enough that it can be interrogated with SNMP. Other active equipment types (e.g., repeaters and bridges that don't speak SNMP) are not used. This constraint:

  1) Helps turn intermittent failures (which are difficult to diagnose) into hard failures.

  2) Often, it allows us to pinpoint the cause of the failure remotely.

- We try to "never backtrack." The network is designed to grow by adding or upgrading equipment. This means that, from time to time, we find ourselves specifying equipment that it "too big" for the current need and thus we are tempted to cut corners. We resist, as we know that we will later spend more time and energy undoing the "temporary" equipment than was saved.

- We use standard "cookbook" network design. This constraint has many benefits:

  1) There is no need for a detailed component-by-component network map. The cookbook is small enough that everyone can know it by heart.

  2) We learn these configurations thoroughly, and can quickly transfer learning from one person or installation to another.

  3) We minimize the amount of inventory that we must stock and that a repair person must take into the field. This means that we are unlikely to be out of stock on a component. Such a stock outage would imply a long down time.

Summary

Large, complex network management systems are outdated. Having one is a sure sign that your network is not designed properly. Instead, your network management system should be "lightweight" and have a minimum of data that you must enter. Other information is derived from this data and from the operating network.

Given that your existing network requires a complex network management system, a necessary part of any solution is to redesign your network according to principles. These principles follow from the service definition and the MTBF and MTTR figures.

The network design now being implemented at the University of Minnesota shows that it is possible to achieve a lightweight network design.

CRAIG FINSETH holds B.S. degrees in Computer Science and Engineering and Philosophy from M.I.T. He is the author of *The Craft of Text Editing*. Since 1987, he has worked for the Minnesota Supercomputer Center and is a "founding" employee of the University of Minnesota's Networking Services department. He has been involved with many network management and computer security projects, including the development of a state-of-the-art network management system designed to meet the needs of large networks. He has been involved with all aspects of designing, operating, and managing the University's 10,000+ node network. He started and is currently overseeing the final stages of a two-year reconstruction of the University's backbone network.

INET '92: The Start of Something Big

by Larry Press, California State University

Introduction

I have had a Net account since 1974, and during most of that time, it was fun, but of somewhat limited value. However, in the last two or three years the Internet has become an indispensable part of my professional life; it is increasingly where I work and where I meet with colleagues. This change is due to the explosive growth of the Internet from the four-node ARPANET in 1969 to an estimated 727,000 hosts in January, 1992 [4], [6].

INET '92

I am not alone in moving my office to the Internet cyberspace, and in recognition of this growing phenomenon, the newly formed *Internet Society* held its first open conference, *INET '92,* in Kobe, Japan, June 15–18, 1992. I attended that meeting, and will summarize it here.

Internet Society

Vint Cerf, Internet Society president, called the meeting "historic," and I am inclined to agree. This was the first open meeting of what I believe will become an important global organization. The *Internet Architecture Board* (IAB), which has overseen the development and standardization of Internet technology, was merged into the Society.

The Internet Society Board of Trustees has agreed to work toward establishing a cooperative relationship with the *International Telecommunications Union.* They hope to effect a link between the IAB and the *International Telegraph and Telephone Consultative Committee* (CCITT). Their hope is that the de facto Internet standards may gain formal acceptance in the international community.

Growth

Cerf predicts rapid growth for both the Society and the Internet. His predictions for the Internet in the year 2000 are shown below. Society membership is open to users, networking professionals, and organizations. If users join in large numbers, it might evolve into something like the Internet's answer to the AAA.

Number of:		*Year:*	
	1992	2000	2000*
Nets	10^4	10^6	10^8
Routers	10^4	$3x10^6$	10^8
Hosts	10^6	10^8	10^{10}
Service Providers	10^{2-3}	10^{3-4}	10^{3-4}
Users	$5x10^6$	10^9	10^9

Internet Society president Vint Cerf projected continued rapid growth of the Internet. The first column of predictions for 2000 is for conventional desktop and portable systems. The second column includes ubiquitous embedded computers such as those described by Mark Weiser in [7].

The conference was headed by Hideo Aiso, Keio University. Larry Landweber, University of Wisconsin, was the Internet Society Liaison, and the program chairman was Haruhuisa Ishida of the University of Tokyo. Professor Ishida fielded a four track program.

Tracks

The first track consisted of *Status Reports on Regional Networks,* covering Africa and the Middle East, Latin America and the Caribbean, Asia and the Pacific, Eastern Europe, Europe, North America, and Japan. These sessions gave one a sense of the global reach of the Internet, and of its heterogeneity.

The second track was on *Network Policy*. This included discussion of difficult issues such as the technical and administrative feasibility of sustaining explosive growth, privacy and security (is privacy guaranteed, or can someone, perhaps in the government, see your mail), and appropriate use (the Internet was established for research and education, how should commercial traffic be treated). The problems in these areas are exacerbated by the global nature of the Internet, which means coexisting with many national governments. As many speakers mentioned, the First Amendment is a local ordinance. Some technical proposals with policy implications were also presented. One was a description of Privacy Enhanced Mail, which provides optional encryption along with assurance that mail came from the indicated originator and was not altered in transit.

The third track was on *Applications*. There were sessions on the role of libraries, networks and social change, entry-level systems appropriate for lesser-developed nations, network management, computer-supported cooperative work, and distance education. While today's impacts and applications are far reaching, they are only a taste of what we will see on tomorrow's high-penetration, high-speed, commercialized Internet.

The fourth track was on *Technology,* with discussions of high-speed ATM switching, next-generation technology, network operations and measurement, addressing and flow control, and multimedia. Notable presentations included a *Virtual Internet Protocol* for portable nodes (significantly it was made by an engineer from Sony), and talk (no paper) by Dave Farber of the University of Pennsylvania, in which he pointed out that a gigabit network is so fast that it might be more appropriate for programmers to view it as a directly-addressable object store in a massively distributed machine than as an input/output "device."

Developing Nations

While the Internet is growing rapidly, it is highly concentrated in developed nations (see Figure 1 on the next page). Since the communication infrastructure in developing nations is typically very poor, Internet connectivity for e-mail and data access is of great marginal value, [1], [2], [5]. As such, a major portion of the meeting was devoted to low-cost, appropriate technology networks for developing nations. (Of course Internet access is unevenly distributed in developed nations. For example, some Cuban and Malaysian school children enjoy better connectivity than those in Los Angeles inner-city schools.)

The conference was preceded by a two-day workshop and tutorial for developing nations. These were organized by Enzo Pullati of the UN Development Program (UNDP), Steve Ruth of George Mason University, and Stefano Trumpy of The National Research Council of Italy. Approximately 60 networkers attended, with financial support of the UNDP and the conference sponsors. There were presentations on the following topics:

- Connectivity and phone company options,
  Joe Choy, choy@ncar.ucar.edu
- The necessity for the networking community to be responsive to the needs of government agencies in developing nations,
  George Sadowsky, sadowsky@nyu.edu
- Communication infrastructure, hardware, software, people, and applications for appropriate technology networks,
  Larry Press, lpress@venera.isi.edu

continued on next page

INET '92: The Start of Something Big *(continued)*

- The evolution of the Southern African regional network from *FidoNet* to *uucp* to IP connectivity, Mike Lawrie, ccml@hippo.ru.ac.za
- The evolution of the Western European network from beginning 300 bps, dial-up uucp connections to its present state, Daniel Karrenberg, karrenberg@ripe.net
- PC-based routers with a live demonstration, Ted Hope, hope@huracan.cr
- *SateLife,* a low-orbit satellite-based system for delivery of medical data and e-mail, Jonathan Metzger, pnsatellife@igc.org
- Internet Services, Art St. George, stgeorge@unmb.bitnet

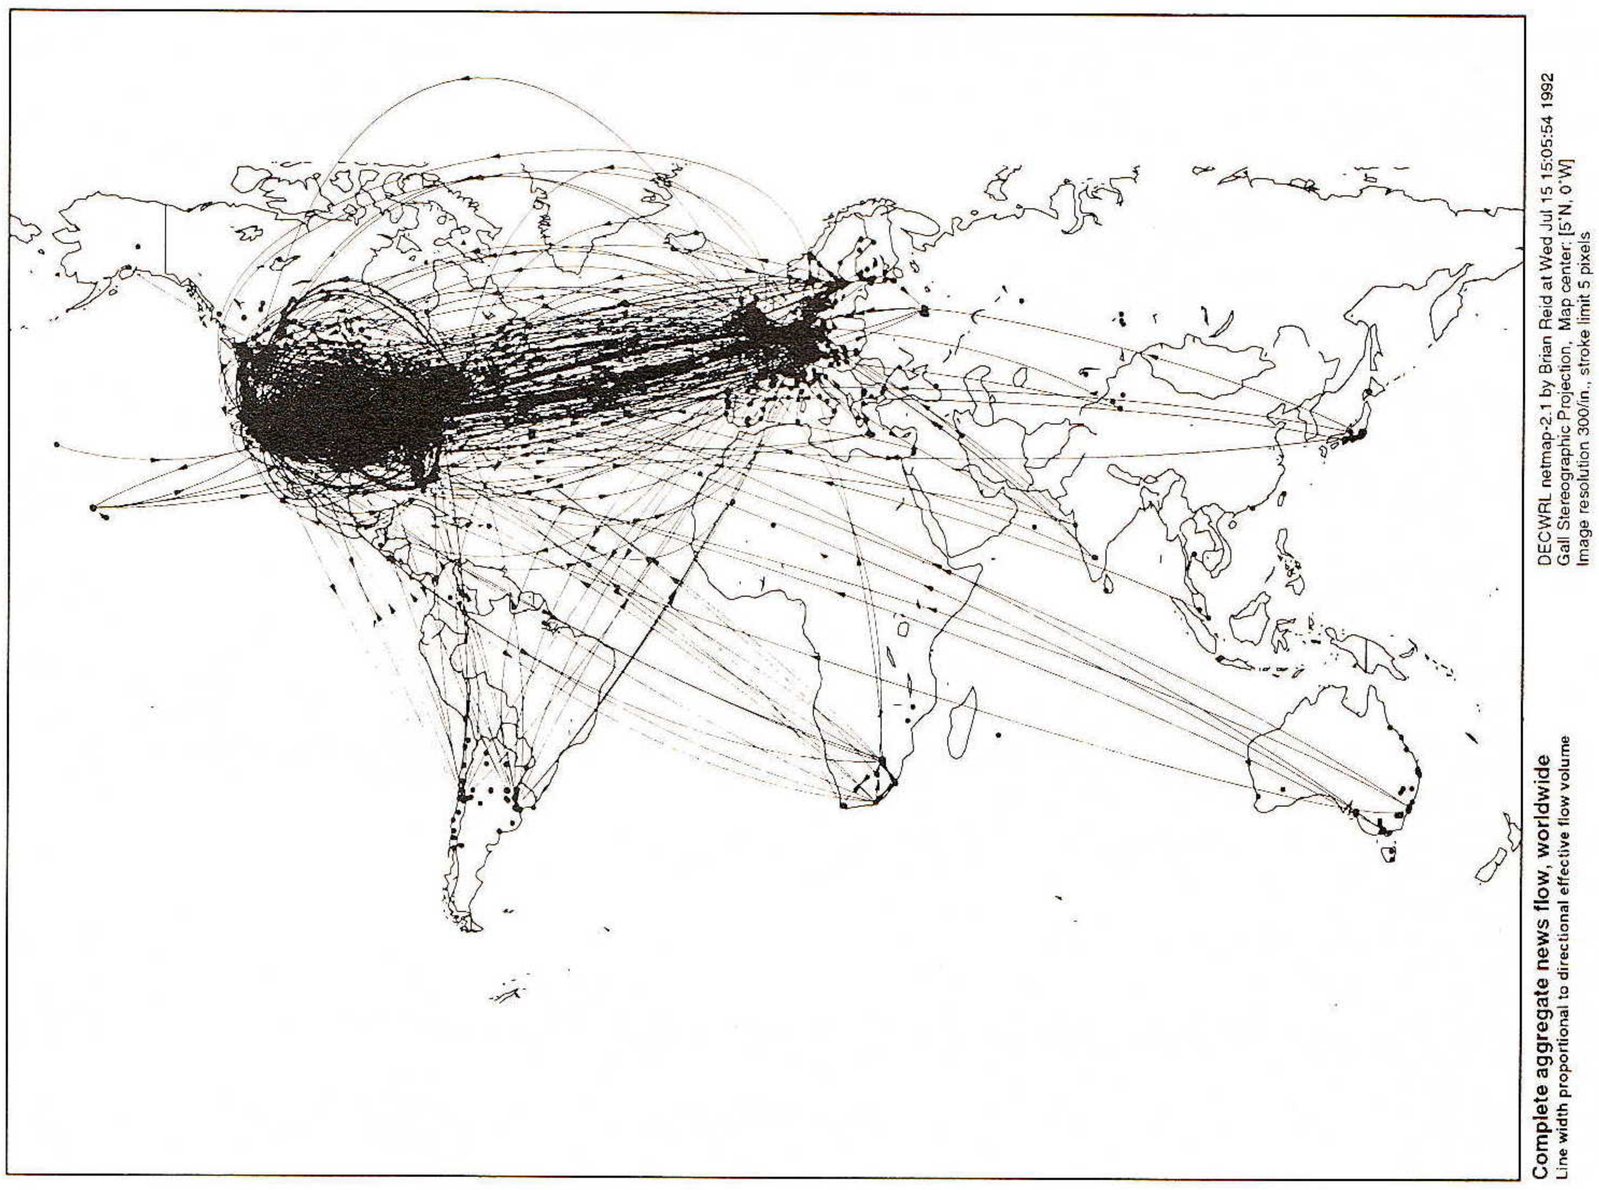

USENET and other international networks are concentrated in North America and Western Europe. In this map, prepared by Brian Reid, line width is proportional to flow volume. Landweber [3] states that 129 of the 236 world political entities lack even e-mail connectivity.

Figure 1: USENET News Flow

In addition to the workshop and tutorial, networks in developing nations were featured in conference sessions on regional networks. The workshop attendees also established a mailing list and an archive for documents on low-cost networking, and an anthology of papers is under preparation. (The archive is in the directory global_nets, available via anonymous FTP at dhvx20.csudh.edu).

Conclusion

INET '92 was an outstanding conference. The papers and presentations were of high quality; the arrangements were excellent (we had terrific food and a room full of Internet-connected workstations—what more could an Internaut ask for); and there was a strong sense of shared enthusiasm and mission. As the song says, it was the start of something big.

For further information about The Internet Society contact:

Internet Society Secretariat
1895 Preston White Drive
Suite 100
Reston, VA 22091
USA
Tel: +1 703 620 8990
Fax: +1 703 620 0913
E-mail: isoc@nri.reston.va.us

References

[1] Dyson, Esther, "Eastern Europe Trip Report, Release 1.0," May 31, 1990, EDventure Holdings, New York, pp 1–30.

[2] Ezigbalike, I. Chukwudozie and Ochuodho, Shem J., "E-Mail for Developing Countries—What They Never Tell You About It," shem@minster.york.ac.uk.

[3] Landweber, L. H., "International Connectivity, Version 5.1," April 20, 1992, lhl@cs.wisc.edu.

[4] Lottor, Mark, "Internet Growth (1981–1991)," RFC 1296, January, 1992.

[5] Press, L., "Relcom: An Appropriate Technology Network," Proceedings of INET '92: The International Networking Conference, Kobe, Japan, June, 1992.

[6] Solensky, Frank, "The Growing Internet," *ConneXions,* Volume 6, No. 5, May 1992, pp 46–48.

[7] Weiser, Mark, "The Computer for the 21st Century," *Scientific American,* September, 1991, pp 94–104.

[8] Goldstein, S. & Michau, C., "Convergence of European and North American Research and Academic Networking," *ConneXions,* Volume 5, No. 4, April 1991.

[9] Stockman, B., "Current Status on Networking in Europe," *ConneXions,* Volume 5, No. 7, July 1991.

[10] Crowcroft, J., "International Internetworking," *ConneXions,* Volume 2, No. 4, April 1988.

[11] "Réseaux Associés pour la Recherche Européenne (RARE)," *ConneXions,* Volume 6, No. 1, January 1992.

[12] Mike Lawrie, "Research and Academic Networking in South Africa," *ConneXions,* Volume 5, No. 8, August 1991.

[13] Steve Neighorn, Randy Bush, and Jeff Beadles, "Profile: RAINet," *ConneXions,* Volume 6, No. 5, May 1992.

[14] Geoff Huston, "Profile: AARNet—The Australian Academic and Research Network," *ConneXions,* Volume 4, No. 3, March 1990.

[15] Mats Brunel, "Profile: NORDUnet," *ConneXions,* Volume 4, No. 11, November 1990.

[16] Mark Bennett, "Electronic Mail in Zambia," *ConneXions,* Volume 6, No. 9, September 1992.

[Ed.: See also "GNET: an Archive and Electronic Journal," page 30 and the INET '93 Call for Participation, page 32.]

LARRY PRESS is professor of Computer Information Systems at California State University at Dominguez Hills in Los Angeles, and a Contributing Editor to the *Communications of the ACM.* He is interested in many topics (too many for his own good), including networking in lesser developed nations and regions (like South Central Los Angles). He attended INET '92, where he participated in the developing nations workshop, and presented a paper on Relcom, a Russian network.

Book Reviews

The Whole Internet User's Guide and Catalog, by Ed Krol, O'Reilly and Associates, ISBN 1-56592-025-2, 1992.

Ed Krol is best known for RFC 1118, "The Hitchhiker's Guide to the Internet." O'Reilly has said "that this is will be the biggest trade title O'Reilly has ever published" and they "believe it will be one of the more popular holiday gift-giving books for computer enthusiasts." [Ed.: The first print run of 15,000 copies has already sold out!]

Audience

In the introduction, Krol says that this book is for professionals, but *not* computer professionals. I think they are missing a potentially large market by this statement. I know many computer professionals especially in the microcomputer arena who do not know anything about the Internet and IP networking, but need to learn in the next few months. Also, Krol said, "if I did my job in writing this book, the rest of what you need you should learn along the way." Let's see if he did.

The introduction of the book says that you do not have to be a UNIX user to use this book, but almost all of the examples in the book are UNIX based. The mail section depends heavily on Berkeley mail and the news section on *nn*. When I read these sections, I feel like I am reading "Introduction to *Berkeley Mail*" and "Introduction to *nn*" chapters. My personal feeling is that the examples will cause more confusion to the novice Internet user who is not using these packages than it does good. If the user is using these packages, then these chapters will be excellent.

UNIX bias

Related to this, the book seems to have a definite UNIX bias, the UNIX information tends to be accurate and complete for the most part. I found several errors, over-generalizations, and oversights in the non-UNIX material. O'Reilly is considering making the book more oriented to the non-UNIX person in the second edition. However, I feel the first edition is lacking in its coverage of non-UNIX systems.

A few places in the book make me feel like I am reading about the perfect Internet. Krol breezes over numeric IP addresses, which is a good thing since many books go into it much too deeply (what end user actually cares what a Class B network is?), but then he goes on to say "Don't worry; you don't need to remember numbers like these to use the network." I wish things were this way, but I know they are not. Many old, yucky TCP/IP packages do not support only the host file and not *Domain Name Service* (DNS). I know this software is out there in use because any time I mention my FTP site without also giving the numeric address I get several mail messages begging me for it from people without DNS. Also, many sites (e.g., my girlfriend's university) are not registered in DNS and hence numeric IP addresses are the only way to get there. Another example is "*Gopher* does not allow you to access anything you couldn't get to already." While I agree this is the way it should be, it is not. I know of quite a bit of information that is only available via *Gopher* (e.g., UNT *NewMan* newsletter—yes, I am as guilty as the next person).

Analogies

Stylistically, the book depends heavily on analogies. Therefore, if you learn well with analogies, this book will be good for you. The book contains many funny inside jokes if you know where to look especially for you PDQ Bach fans.

Strengths

One of the strengths of the book in my opinion is when it talks about social issues on the Net. I cheered when I read Krol's comments on USENET censorship. The etiquette sections were also good.

However, the book suffers from inconsistency in spots. Why is USENET *News* mentioned as being available under *WorldWideWeb* and not *Gopher*? (Note: USENET *News* in *Gopher* has been around since February or March.) Why is *Knowbot* mentioned and *Netfind,* which I think is an infinitely more useful tool, not? Why is something as basic as the incompatibles between new and old talk not mentioned, but fancy FTP "get" commands involving pipes mentioned? I was not even aware of these "get" commands until I read the book.

The *Gopher, WWW, WAIS,* and resource catalog sections were a bold move. This information is a fast moving target. Fortunately, while I see minor spots that are dated in these sections, the chapters are still fairly current and useful from the end users point of view. The section on dealing with *WAIS* searches that have gone awry is the best discussion on this topic I have seen. Actually, the *WAIS* and *WWW* chapters are among the best discussions of this material that I have seen.

Updates

As a keeper of an Internet resource catalog, I know that resource information goes out of date fast. This somewhat worries me about having a resource catalog in print. O'Reilly has said that they will be updating the book frequently, but the question remains how many people will buy the updates at $24.95. On a more positive note, even I, with an incomplete version of the catalog section, found useful resources that I have not seen before. I think users will enjoy this section and find it useful as long as the information does not become obsolete. There is a vague note in the information sheet about the Resource Catalog being available on the Internet at a later date, but I do not know any of the details. If implemented correctly, it would vanquish my obsolete information fear.

While the "Dealing with Problems" section has much good information, parts of it will make network administrators lose sleep. The book talks about fixing Ethernets, but the note saying that you (the user) may want to check with your network administrator before doing anything comes at the end of the chapter. I fear that many users will read up to the repairing part and not get to the note before attempting to do something.

Zen

To compare this book with the first/free edition of "Zen and the Art of Internet" I recently reviewed (I have not yet seen the second edition), I would have to say that this book is much longer and more complete. The other side to this is that "Zen" explains topics in a clear, compact format whereas Krol's is much more verbose. The most noticeable organizational difference is that "Zen" does not have a resource catalog or chapters on some of the newer Internet utilities: *WAIS, WWW,* or *Gopher.*

Lukewarm

As you have guessed by now, I am lukewarm on this book. I would suggest buying this book if you want the unique and useful resource guide, if you need information on *WAIS, WWW,* or *Gopher,* or if you are using *Berkeley Mail* or *nn* on a UNIX box. Otherwise, give the The *Whole Internet User's Guide and Catalog* the same looking over that you give "Zen and the Art of the Internet" and some of the soon-to-be released books on the same topic. I would also advise you computer professionals out there to ignore the audience section of the book and take a look at it anyway if you need to get up to speed on the Internet. Finally, the second edition of this book should be worth reviewing again if it fixes the problems of the first edition.

—Billy Barron, University of North Texas (billy@unt.edu)

continued on next page

Book Reviews *(continued)*

XTP: The Xpress Transfer Protocol by W. Timothy Strayer, Bert J. Dempsey, and Alfred C. Weaver, Addison-Wesley, 1992 (ISBN 0-201-56351-7).

Well-written

A book about a single protocol is a daunting idea (imagine finding enough to say about TCP to fill 200 pages), but these authors pull it off. This is a well-written book, but I'm not sure XTP is worthy of this effort.

Organization

The book has two parts. The first part, which consists of the first three chapters, plus introductory parts of some later chapters, present the motivation for XTP by surveying key problems that XTP was designed to address, such as high throughput and fast connection establishment. The second part is a detailed description of the workings of XTP. I have no concerns about the second part, and in some places, found it intriguing (for example, regarding some of the ideas about multicasting). But the first part of the book disturbed me.

The goal of the first part of the book is to explain why a new protocol like XTP is needed. While reading these early chapters, I was reminded of an insightful comment by a friend about how to read theology texts: the trick is to notice what parts of the Bible the writer isn't mentioning, because they're often the parts that hurt the theory. And there's a lot of well-known work that hurts the XTP cause that isn't mentioned. While these omissions may not be intentional, they make these chapters less credible.

Lightweight protocol

The fundamental idea behind XTP is that to go fast, one needs to develop a so-called "lightweight" protocol (a term the XTP designers invented). The XTP folks define a lightweight protocol as one that is easy to parse, provides a useful range of transport services, requires a minimum amount of CPU support, and can be largely implemented in hardware. It is hard to argue with the idea of such protocols (except regarding implementing them in hardware, about which there are bitter arguments—some implementors think bad interface hardware is the source of many performance problems). But there's a strong argument that existing protocols, such as TCP/IP and TP4, actually meet this definition, and the book avoids mentioning work that suggests TCP/IP or TP4/CLNP can run just as fast or faster than XTP.

Comparison with other work

For example, much of the motivation for XTP comes from some performance measurements done in the late 1980s showing that some commercial TCP/IP and TP4 implementations were very slow. These studies are cited in Chapter 1 as showing the need for a higher-performance protocol. But the book does not cite the studies done after the performance measurements that showed that the reasons for poor performance were not the protocols but rather how they were implemented (the performance measurements were taken with implementations that are now viewed as the Trabants of the networking field) and the properties of the interface chips used. The cost of processing a TCP segment or forwarding an IP datagram in the best implementations is now measured in *tens* of instructions. It to show a burning performance requirement to replace a protocol that only consumes tens of instructions per packet. (For more details on TCP performance, see the notes for Van Jacobson's tutorial at SIGCOMM '90, or see his note in the April '90 issue of *Computer Communication Review,* the article by Clark et. al. in the June '89 issue of *IEEE Communications Magazine,* and Borman's articles on the 800 Mbit/s Cray TCP/IP implementation in the January '91 issues of *Computer Communication Review*).

Omissions

There were a number of more minor omissions in Chapter 3. The chapter states that computing accurate round-trip times is hard and cites a 1986 study, but never mentions the algorithms of Karn (paper at SIGCOMM '87) and Jacobson (paper at SIGCOMM '88), which are widely credited with solving the problems identified in the 1986 study. It characterizes the DEC-bit and slow-start congestion control schemes as "largely reactive and fairly crude," without supporting citations, when others have argued that these algorithms are close to optimal in their behavior over networks with varying loads. It also states that "in general, protocols require more processing per packet at the receiver than at the sender," but never mentions that Jacobson's work on header prediction strongly suggests that processing costs are actually higher for the sender. Finally, Chapter 3 states that TCP uses "go-back-n" retransmission. In fact, TCP's retransmission is more like selective retransmission than go-back-n.

If you want to learn more about XTP, go ahead and buy the book. It is well-written. But do keep in mind that there are other points of view.

—Craig Partridge, BBN Systems and Technologies

The Internet Companion: A Beginner's Guide to Global Networking, by Tracy LaQuey and Jeanne C. Ryer with foreword by Senator Al Gore is the latest book on the Internet to hit the market. The cost is a low $10.95, which is half the cost of any of the other Internet books currently on the market. The ISBN is 0-201-62224-6.

The basics

The title says it is a beginner's guide and *The Internet Companion* is exactly that. If you are hoping to become a power user, then buy a copy of *Zen and the Art of the Internet* or *The Whole Internet User's Guide and Catalog* instead. *The Internet Companion* is directed to the beginning novice especially one who is either not currently connected or not sure why they should use the Internet. It even covers beginning topics such as what files and accounts are.

Anecdotes and examples

A good portion of the book is spent on selling the Internet. This is largely done through the use of anecdotes which range from the cultural "Enough of White Man's ASCII" to the political "Serious Games" and silly "Elvis Sighted on the Internet." They provide quick glances into all aspects of Internet life and should excite some potential users.

One of the principles used in writing this book was to only show examples when these would help and not confuse the user due to system differences. Thus you will find examples of Telnet and FTP commands, which are pretty universal, but not of e-mail or USENET because there are dozens of very different mail and news packages. Since I have complained about the dependence on a single package in other Internet books, I found this a refreshing change.

A chapter called "Getting Connected" is directed at the user who is in need of an individual or small business Internet connection. It explains all the major issues involved with getting connected. Then it goes on to list the network providers that offer this type of service. After this chapter, the book closes with a very complete bibliography.

Recommended

In conclusion, if you are already a competent Internet user, do not buy this book. If you are a beginner or looking to get a home Internet connection, then this is the book for you. If you know people who could benefit from the Internet and you want to sell them on the idea, get them to read this book. Finally, *The Internet Companion* should reach many people due to its low price.

—Billy Barron, University of North Texas (billy@unt.edu)

GNET: an Archive and Electronic Journal

Toward a Truly Global Network

Introduction

Computer-mediated communication networks are proliferating and growing rapidly, yet they are not truly global—they are concentrated in affluent parts of North America, Western Europe, and parts of Asia.

GNET is an archive/journal for documents pertaining to the effort of bringing the net to lesser-developed nations and the poorer parts of developed nations. (Net access is better in many "third world" schools than in South-Central Los Angeles). GNET consists of two parts, an archive directory and a moderated discussion list.

Documents

GNET documents are available by anonymous FTP from the directory `global_net` at `dhvx20.csudh.edu`. (address 155.135.1.1). To conserve bandwidth, the archive contains an abstract of each document as well as the full document. (Those without FTP access can contact me for instructions on mail-based file retrieval).

Mailing list

In addition to the archive directory, there is a moderated GNET discussion list. This list is limited to discussion of the documents in the archive. It is hoped that document authors will follow this discussion, and update their documents accordingly. If this happens, the archive will become a dynamic journal. Monthly mailings will list new papers added to the archive.

Topics

We wish broad participation, with papers from nuts-and-bolts to the visionary. Suitable topics include, but are not restricted to:

- Descriptions of networks and projects
- Low-cost, appropriate-technology networks
- Satellite and terrestrial packet radio
- Communication protocols
- Connection options
- Host and user hardware and software
- The current state of global networking
- Current applications
- Proposed applications
- Education using the global net
- Social, political or spiritual impact
- Economic and environmental impact
- Politics and funding
- Political implications of a global network
- Free speech, security and privacy
- Directories and lists of people and resources

More information

To submit a document to the archive or subscribe to the moderated discussion list, send a message to:

`gnet_request@dhvx20.csudh.edu`.

—Larry Press, California State University

IVS: Software for Videoconferencing on the Internet

Goals

The *INRIA Videoconferencing System* (IVS) is designed with the explicit goal to demonstrate that it is possible today to use a standard workstation for video-/audio conferences. Thus, in order to conduct a videoconference with IVS, only minimal hardware upgrades are required to a machine commonly found on the desk of an engineer (a Videocamera and a Videocard). A further design goal of IVS is to allow more than two persons to conduct a videoconference using IP multicast extensions. Thus, with IVS, a group of people can participate in a conference at the same time.

IVS allows use of standard Internet technology for transmission of video/audio-data. This is achieved by implementing a software version of a H.261 codec (video codec for audiovisual services at p*64 Kbps) in IVS. Furthermore, where conventional H.261 hardware codecs require leased lines or switched circuits for data-transmission, the H.261 software codec of IVS uses standard UDP datagrams.

Development

This system has already been tested in many countries including France, Germany, Sweden, the Netherlands and the USA. The following also contributed to the development of IVS:

David Berry (Videopix grabbing improvement)
`David.Berry@eng.sun.com`

Jack Jansen (ADPCM 32Kbps audio codec)
`Jack.Jansen@cwi.nl`

Guido van Rossum (SGI Indigo compatibility)
`Guido.van.Rossum@cwi.nl`

Winston Dang (Speed improvements)
`wkd@uhunix.uhcc.hawaii.edu`

Technical data

- *System requirements:* Sun SPARCstation or SGI Indigo stations, Video grabber (VideoPix Card for Sun workstations), Video Camera, The X Window System, and Internet access.
- *FTP address:* IVS Version 1.10 is available for anonymous FTP:
  Host: `avahi.inria.fr`
  Directory: `/pub/videoconference/ivs.tar.Z`.

—Thierry Turletti (`turletti@jerry.inria.fr`)
Project RODEO—INRIA Sophia-Antipolis—France

Document Update

New names

The SMP drafts mentioned in the October issue of *ConneXions* have recently been renamed since they are now being edited by the IETF's SNMP Version-2 (SNMPv2) working-group. They are now called:

```
draft-ietf-snmpv2-coex-01.txt
draft-ietf-snmpv2-intro-01.txt
draft-ietf-snmpv2-m2m-01.txt
draft-ietf-snmpv2-mib-01.txt
draft-ietf-snmpv2-proto-01.txt
draft-ietf-snmpv2-smi-01.txt
draft-ietf-snmpv2-tc-01.txt
draft-ietf-snmpv2-tm-01.txt
```

Location

You should be able to find them in all the usual locations for Internet Drafts. (We retrieved the above list from the `~/internet-drafts` directory at host `nnsc.nsf.net`).

Call for Participation

Following the very successful INET '92, the *INET '93 International Networking Conference* will be held on 17–20 August 1993 in San Francisco, California. Focusing on worldwide issues of research and academic networking, the goal of INET '93 is to bring together individuals from university, industry and government who are involved with planning, developing, implementing, managing and funding national, regional and international research, academic, and commercial networks. The conference is hosted by the Internet Society.

Topics

The official language of the conference is English. The conference agenda will include plans and status reports on research and academic networks throughout the world. In addition, possible topics for conference sessions include but are not limited to the following:

Network Technology—Advances in the Network Technology Base:

- Progress toward international open network protocols
- Security, management and authentication in managing networks
- Transmission, routing, and transport technologies
- Technologies of the '90s and the 21st century
- Very high speed networks

Network Engineering—Building the Global Infrastructure:

- Application of network technology to provide networking services
- Interoperability among existing national/international networks
- Network management systems and methods
- Reliability and performance engineering
- Issues related to scaling

INET '93

Enabling Technologies for Distributed Applications:

- Collaboration technologies
- Multimedia issues
- Networked information retrieval
- Mail and directory services
- Workstation teleconferencing
- Computer supported collaborated work
- Interoperability of application services

Support for International Communities of Interest:

- Support of international collaboration
- Access to scientific papers and data across national boundaries
- Supercomputing
- High energy physics, atmospheric modeling, and other scientific applications
- Education/distance learning
- Medical research and clinical applications
- Libraries

Work and play in Cyberspace—How networks are changing the social nature of work and play:

- Networking and the arts
- High payoff application areas to support national and international development

Regional Issues—Networking Around the Globe:

- Unique regional issues and approaches such as multilingual and national character set accommodation

- Asia-Pacific Rim
- Eastern Europe
- Europe
- Former Soviet Republics
- Latin America
- North America
- Africa
- Special Issues for the Third World

Policy Issues—Governance, Management, and Financing of International Networks:

- Globalization of services
- Commercialization, privatization and public access
- Coordination of international resources
- Copyright and intellectual property rights
- Appropriate use and speech restrictions
- International security policy
- Privacy and data protection
- Telecommunications policy

INET '93 and INTEROP

The conference will be held immediately preceding INTEROP 93 Fall, the leading conference and trade show for Internet technologies. This will make possible attending both events as well as tutorials given as part of INTEROP 93 Fall.

The conference will be held in one of the most beautiful cities in the world. Social events will be arranged during the conference to take advantage of the unique environment San Francisco provides. In addition, assistance will be provided to attendees wishing to see more of this unique area.

Submissions

Please submit 6 copies (in English) of double-spaced typed manuscript (maximum of 20 pages) with an abstract to:

USRA
Attn.: INET'93
625 Ellis Street, Suite 205
Mountain View, CA 94043
Tel.: +1 415 390-0317
Fax: +1 415 390-0318

You may also submit an electronic (ASCII, please) version of your paper by e-mail to:

`Submission@inet93.stanford.edu`

Important dates

January 10, 1993	Abstracts due
March 1, 1993:	Manuscript due
May 1, 1993:	Notification of acceptance to authors
June 10, 1993:	Camera-ready papers due

Workshop for Developing Countries

A workshop designed to assist developing countries in their installation and use of networking technology is being organized and will take place during the week before the conference in the San Francisco Bay Area.

More information

To be added to the conference mailing list or for other requests, send mail, fax to the postal address above or e-mail to:

`Request@inet93.stanford.edu`

Call for Papers

The 1993 *International Conference on Network Protocols* (ICNP-93) will be held at the ANA Hotel, San Francisco, California, USA, October 19–22, 1993.

Topics

ICNP-93 is sponsored by the IEEE Computer Society Technical Committee on Distributed Processing. Original technical papers addressing the following topics of interest are solicited for presentation at the conference and publication in the conference proceedings:

- Network Architectures
- Switching Protocols
- Routing Protocols
- Flow & Congestion Control
- High-Speed Networks
- Real-Time Protocols
- Network Security
- Name Servers & Directories
- Protocol Conversion
- Broadcast Systems
- Distributed Operating Systems
- System Support and Interfaces
- Protocol Design Methodology
- Protocol Verification
- Protocol Testing and Debugging
- Protocol Implementation

Submissions

Authors are requested to submit six copies (in English) of their double-spaced typed manuscript (maximum of 25 pages) with an abstract to the program chair by March 1, 1993. Papers will be considered for publication in the new journal *IEEE/ACM Transactions on Networking*. Submit papers to:

Prof. Mohamed G. Gouda, Program Chair
Department of Computer Science
University of Texas
Austin, Texas 78712
USA
Tel: 512-471-9532
E-mail: gouda@cs.utexas.edu

Tutorials

The day before the conference (October 19) will be devoted to tutorials for providing the background on the conference. Send tutorial proposals by March 1, 1993 to:

Dr. M. Umit Uyar, Tutorials Chair
AT&T Bell Labs, Room 3D-501A
Crawfords Corner Road
Holmdel, New Jersey 07733
USA
E-mail: umit@honet5.att.com

Important dates

Papers and proposals due:	March 1, 1993
Acceptance letters sent:	June 15, 1993
Camera ready copies due:	August 1, 1993

More information

For further information, please contact:

Prof. Ming T. (Mike) Liu, General Chair
Department of Computer and Information Science
The Ohio State University
2036 Nell Ave.
Columbus, Ohio 43210
USA
E-mail: liu@cis.ohio-state.edu

Electronic Access to ITU (e.g., CCITT) Documents

Teledoc

The *International Telecommunication Union* (ITU), a United Nations Agency based in Geneva, Switzerland, has announced a new electronic document distribution service called *Teledoc.* The ITU consists of five permanent organs including the General Secretariat, the International Frequency Registration Board (IFRB), the International Radio Consultative Committee (CCIR), the International Telegraph and Telephone Consultative Committee (CCITT) and the Telecommunications Development Bureau (BDT).

Access methods

The *Teledoc* service makes available public ITU documents in a database called the "ITU Document Store." The ITU Document Store organizes ITU documents into hierarchical groups: each group can contain additional groups and/or documents. Remote access to the ITU Document Store is planned via:

- Electronic mail (auto-answering mailbox)
- Interactive VT interface (planned for early 1993)
- Internet FTP (planned for early 1993)

TAM

The first available interface is the *Teledoc Auto-Answering Mailbox* (TAM), an X.400-based document server. Mail messages can be sent to:

X.400: `S=teledoc; P=itu; A=arcom; C=ch`

or

Internet: `teledoc@itu.arcom.ch`

Commands to the TAM must be placed in the mail message body (not in the subject field). The commands are simple. For example:

```
HELP
LIST CCITT
LIST CCITT/REC
```

will send the TAM help file and a list of the contents of the CCITT and CCITT Recommendations group (only lists of CCITT Recommendations and some summaries are available as of October 1992). The HELP file describes how to retrieve individual documents.

More information

For additional information about *Teledoc,* please contact:

Robert Shaw
Teledoc Project Coordinator
Information Services Department
International Telecommunication Union
Place des Nations
1211 Geneva,
Switzerland
Voice: +41 22 730 5338/5554
Fax: +41 22 730 5337
Internet: `shaw@itu.arcom.ch`
X.400: `G=robert;S=shaw;P=itu;A=arcom;C=ch`

[Ed.: We will have more information about this service in a future issue.]

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1